

AFIT/GCS/ENG/97D-18

A PERFORMANCE ANALYSIS OF THE
FAUGERAS COLOR SPACE AS A COMPONENT
OF COLOR HISTOGRAM-BASED IMAGE RETRIEVAL

THESIS
Chad A. Vander Meer
2nd Lt, USAF

AFIT/GCS/ENG/97D-18

19980127 073

| DTIC QUALITY INSPECTED 3

Approved for public release; distribution unlimited

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U. S. Government.

A PERFORMANCE ANALYSIS OF THE
FAUGERAS COLOR SPACE AS A COMPONENT
OF COLOR HISTOGRAM-BASED IMAGE RETRIEVAL

THESIS

Presented to the Faculty of the School of Engineering

of the Air Force Institute of Technology

Air University

In Partial Fulfillment of the

Requirements for the Degree of

Master of Science in Computer Engineering

Chad A. Vander Meer, B.S.CompE.

2nd Lt, USAF

December, 1997

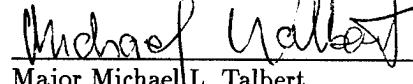
Approved for public release; distribution unlimited

A PERFORMANCE ANALYSIS OF THE FAUGERAS COLOR SPACE AS A
COMPONENT OF COLOR HISTOGRAM-BASED IMAGE RETRIEVAL

Chad A. Vander Meer, B.S.CompE

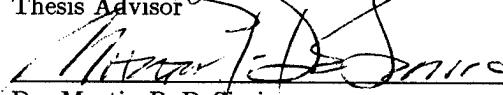
2nd Lt, USAF

Presented to the Faculty of the Graduate School of Engineering
of the Air Force Institute of Technology
In Partial Fulfillment of the
Requirements for the Degree of
Master of Science in Computer Engineering



Major Michael L. Talbert
Thesis Advisor

25 Nov 97
Date



Dr. Martin P. DeSimio
Member, Thesis Committee

25 Nov 97
Date



Dr. Matthew Kabrisky
Member, Thesis Committee

25 Nov 97
Date

Acknowledgements

Many individuals deserve thanks for their contributions to the successful completion of this thesis. Those mentioned below provided an overwhelming amount of guidance and encouragement.

First, I'd like to thank my thesis advisor Maj Michael Talbert. His comments and suggestions not only guided my research, they vastly improved my skills as a writer. Thanks to Dr. Martin DeSimio for the insight gained from our meetings (Although there weren't many meetings, each one was extremely beneficial for keeping my research in perspective). To Dr. Matthew Kabrisky and Captain John Keller, thanks for making this experience so memorable. Your encouragement and enthusiasm allowed me to realize the true enjoyment that can be attained through research. Also, to my sponsor, Major Steve Mateschik at Rome Labs, thanks for allowing me flexibility in exploring the problem domain of content-based image retrieval.

Finally, thanks to all my friends, classmates, and especially my wife Amanda. Although we have only been married for a few months, her strength (and cooking) have helped me weather the bad times more than once.

Chad A. Vander Meer

Table of Contents

	Page
Acknowledgements	iii
List of Figures	vii
List of Tables	viii
Abstract	ix
I. Introduction	1-1
1.1 Overview	1-1
1.2 Problem Statement	1-4
1.3 Scope	1-4
1.4 Assumptions	1-4
1.5 Approach	1-5
1.6 Thesis Organization	1-6
II. Background	2-1
2.1 Introduction	2-1
2.2 Indexing	2-2
2.2.1 Overview	2-2
2.2.2 Considerations	2-3
2.3 Color Histograms	2-5
2.3.1 Overview	2-5
2.3.2 Feature Vectors	2-5
2.3.3 Construction and Use of a Color Histogram	2-7
2.3.4 Methods for Improving Color Histogramming	2-10
2.3.5 Summary	2-16
2.4 Human Visual System	2-16

	Page
2.4.1 Overview	2-16
2.4.2 Image Encoding	2-17
2.4.3 Image Representation	2-18
2.4.4 Image Interpretation	2-19
2.5 Faugeras Color Space	2-20
2.5.1 Introduction	2-20
2.5.2 Assumptions	2-20
2.5.3 Faugeras HVS Model	2-20
2.5.4 Limitations of Faugeras Model	2-23
2.5.5 Benefits of Using Faugeras Color Space	2-24
2.6 Summary	2-24
 III. Methodology	 3-1
3.1 Introduction	3-1
3.2 Setup	3-1
3.3 Computer Histogram Intersection	3-2
3.3.1 Color Space Transformation	3-2
3.3.2 Generation of Feature Vectors	3-3
3.3.3 Similarity Comparison	3-6
3.4 Perceptual Experiment	3-8
3.4.1 Color Imagery Experiment	3-8
3.4.2 Presentation of Data	3-9
3.4.3 Experimental Procedure	3-10
3.4.4 Use of Resulting Data	3-11
3.5 Summary	3-11

	Page
IV. Results	4-1
4.1 Introduction	4-1
4.2 Analysis of Results from Perceptual Experiment	4-1
4.2.1 Bias Due to Content of Images	4-2
4.2.2 Bias Caused by Misinterpretation of Instructions	4-5
4.2.3 Other Possible Biases	4-6
4.2.4 Summary	4-6
4.3 Correlation Results for Computer Histogram Intersection	4-7
4.3.1 Pearson Correlation	4-7
4.3.2 Plane Average Results	4-7
4.3.3 Uniform Feature Vector Results	4-9
4.3.4 Non-Uniform Feature Vector Results	4-11
4.4 Summary	4-12
V. Conclusions	5-1
5.1 Performance of the Faugeras Color Space	5-1
5.1.1 Faugeras Color Space (with CSF)	5-1
5.1.2 Faugeras Color Space (without CSF)	5-2
5.2 Recommendations for Further Work	5-2
5.3 Summary	5-3
Appendix A. Similarity Matrices	A-1
Appendix B. Experiment Matlab M-Files	B-1
Appendix C. Color Histogram Matlab M-Files	C-1
Appendix D. Test Images	D-1
Bibliography	BIB-1
Vita	VITA-1

List of Figures

Figure	Page
1.1. Breakdown of Imagery Use in the United States.	1-2
2.1. Storage of feature vectors with their corresponding image to allow quick comparisons with query images.	2-6
2.2. A model for content-based retrieval.	2-7
2.3. Three-Dimensional representation of the RGB color space.	2-8
2.4. Eight bin partition of a single color space axis.	2-11
2.5. Layout of the Human Eye.	2-17
2.6. Components of the Faugeras Color HVS Model.	2-21
3.1. Vector produced for RGB color space by the Plane Averages method	3-6
3.2. Instructions describing the experiment.	3-9
4.1. Test Images Dominated by Blue	4-3
4.2. Example of Possible Bias Based on Images used in Tutorial	4-4
A.1. Similarity Matrices for Plane Average Feature Vectors Produced from the RGB and HSV Spaces.	A-2
A.2. Similarity Matrices for Plane Average Feature Vectors Produced from the RGB and HSV Spaces.	A-3
A.3. Similarity Matrices for Uniform Feature Vectors Produced from the RGB and HSV Spaces.	A-4
A.4. Similarity Matrices for Uniform Feature Vectors Produced from the Faugeras Spaces.	A-5
A.5. Similarity Matrices for Non-Uniform Feature Vectors Produced from the RGB and HSV Spaces.	A-6
A.6. Similarity Matrices for Non-Uniform Feature Vectors Produced from the Faugeras Spaces.	A-7

List of Tables

Table	Page
3.1. Matrix format for storing similarity results	3-8
4.1. Pearson r Values for Plane Averages Method	4-8
4.2. Confidence that there is a Difference Between RGB and HSV r Values	4-8
4.3. Pearson r Values for Uniform Method	4-9
4.4. Confidence That There is a Distance Between RGB and Faugeras (without CSF) r Values (and therefore differences in performance)	4-10
4.5. Pearson r Values for Non-Uniform Method	4-11
A.1. Similarity Values Obtained from Human Perceptual Experiment.	A-1
A.2. z Values Obtained for Human Perceptual Results.	A-1
A.3. Measured Variances for Human Perceptual Results.	A-1
D.1. Ten Test Images Used for Experiments	D-1
D.2. Images Used for Tutorial	D-2

Abstract

The use of color histograms for image retrieval from databases has been implemented in many variations since the original work of Ballard and Swain. Selecting the appropriate color space for similarity comparisons is an important part of a color histogram technique. This paper serves to introduce and evaluate the performance of a color space developed by O.D. Faugeras through the use of color histograms. Performance is evaluated by correlating the similarity results obtained from various color feature vector techniques (including color histgramming) to those gathered through a human perceptual test. The perceptual test required 36 human subjects to evaluate the similarity of 10 military aircraft images. The same 10 images were also compared via the color feature vector techniques. The results obtained for the *Faugeras* color space are compared against those of the Red, Green, Blue (RGB) and Hue, Saturation, Value (HSV) color spaces. While the correlation results for the Faugeras color space were unexpected and unfavorable, a Pearson correlation coefficient of 0.91 was obtained for the HSV space suggesting that HSV is an excellent color space for judging color image similarity. A discussion of the Faugeras space's performance and future research directions are presented at the conclusion of the paper.

A PERFORMANCE ANALYSIS OF THE FAUGERAS COLOR SPACE AS A COMPONENT OF COLOR HISTOGRAM-BASED IMAGE RETRIEVAL

I. Introduction

This thesis addresses the problem of determining an appropriate color space representation of digital color images when color similarity calculations must be performed. Assessment of color similarity is an important element of military and commercial applications involving content-based image retrieval from databases.

To evaluate the performance of a new color space in this particular problem domain, several different variations of a single color-based image retrieval technique are constructed. Similarity results are obtained from the retrieval technique and then compared against results collected through a human experiment to assess overall performance.

1.1 Overview

Advancements in computer technology have provided the ability to economically store images, sound, and motion video in digital format. In fact, imagery has become an essential part of everyday business. Two examples of institutions where the importance of digital images has evolved are hospitals and commercial image distribution corporations. Hospitals can produce and be required to store as much as fifty Gigabytes of diagnostic images each day (6), while image distribution companies like R.R. Donnelley and Sons estimate an on-line storage capability of 100 Terabytes for future customer image accessibility needs (7). In addition, the United States government and the military in particular store enormous volumes of imagery. In fact, the government and military account for 35.5% of all U.S. imaging (Figure 1.1).

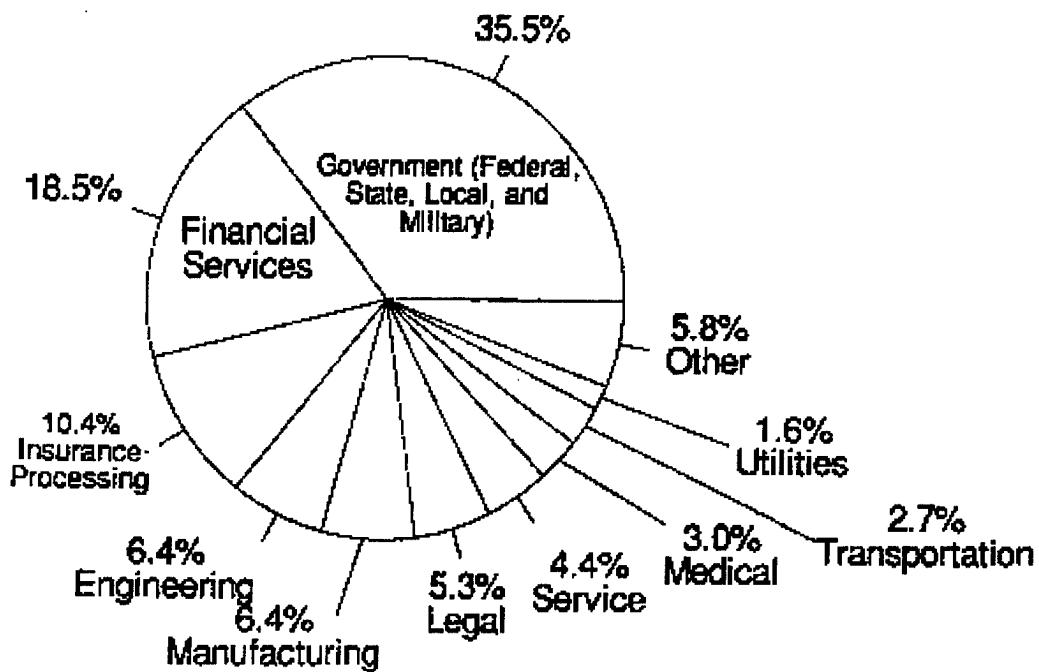


Figure 1.1 Breakdown of Imagery Use in the United States (1).

This rapid accumulation of digital imagery has resulted in a need to further automate the process of searching for and retrieving such files. In the past, multimedia retrieval has been limited to keyword searches based on one person's interpretation of an image's relevant content(2). This technique suffers from a number of flaws. First, images cannot be completely described by a listing of keywords. As the accuracy of the image description increases, the storage requirements also increase. Probably the most expensive requirement of the keyword technique is the time needed by humans to interpret individual images and produce the keyword listing. Finally, the use of keywords prevents the global accessibility of a multimedia archival (2). For example, descriptive words associated with an image are normally from a single language. The use of multiple languages such as English, German, and Japanese is a possibility but adds to the burden of storage requirements. Essentially, the expansion of the Internet has provided a means for international communication requiring the elimination of linguistic barriers.

One solution to the problems facing digital image retrieval is to extract a reduced representation of the image by automated means. Features like color, texture, and spatial information, which are used by humans to assess and remember image content, are manipulated and compared by methods based on human perceptions. While these new content-based retrieval systems help alleviate the problems introduced by keyword searches, the query abilities of such systems are still simplistic (8) and lack the efficiency needed to access massive archives. Improvements in retrieval accuracy (finding images a user wants) have been attained through the extraction of multiple image features. The disadvantage of using multiple features is the increase in retrieval time. One way to help control increases in retrieval time is by improving techniques based on individual features. A robust and efficient method of retrieval is color histogram intersection(9). Although color histograms do not preserve spatial orientations in images, they still provide an important way to judge similarity. Improvement of the color histogram technique has focused in three areas: 1) Similarity metrics (i.e. Euclidean distance, etc.) 2) Color Spaces and 3) Color Space Quantization. The focus of this research was to investigate the value of the Faugeras color space as a meaningful and effective

component of an image retrieval system. The Faugeras color space is based on the physiology of the human visual system and follows from the work of O. D. Faugeras(10).

1.2 Problem Statement

Does the Faugeras color space, when used as a component of color histogramming, help provide better correlation with the human perception of color image similarity than the RGB (Red, Green, Blue) and HSV (Hue, Saturation, Value) color spaces?

1.3 Scope

The scope of this thesis is limited to a comparison of the Faugeras color space with the HSV and RGB color spaces. This will be done in terms of their effect on the correlation between image similarity defined by a Euclidean distance measure and that measured through a human perceptual experiment.

1.4 Assumptions

A few assumptions had to be made in order to implement a non-uniform color quantization method and to consider the similarity comparison process valid. First, the database of images is assumed to be preestablished. Definition of the uniform and non-uniform quantization techniques used in this research rely on prior knowledge of the database color distribution.

Second, a Euclidean distance is assumed to be appropriate as a similarity measure. Although humans do not measure similarity in such a fashion, the Faugeras space's construction was based on a Euclidean vector space. This allows for meaningful, mathematically tractable distance measures that work well for comparison purposes(11).

Finally, color histogram retrieval is most effective when scanning through heterogeneous image collections. For example, in this retrieval domain, a query for images similar to a picture of a taxi

cab will return a collection of images with yellow objects (hopefully some of which are images of taxis). The same similarity query is of little value in a homogeneous database. If all images in the database are of taxis, use of color to discriminate between pictures is highly ineffective since each image's color histogram is nearly identical. Therefore, since color histograms are the featured retrieval technique for this research, the database is assumed to be composed of images with a wide variety of colors and objects.

1.5 Approach

The first task of this research is to produce n -dimensional vectors that can be compared using the color histogram intersection method. A vector is created for each axis of the different color space representations of an image. Three different types of vectors are constructed. They are: 1) a color histogram based on uniform quantization of each image's color distribution 2) a color histogram based on non-uniform quantization of each image's color distribution 3) a collection of three values representing the average of each color space plane. Using a Euclidean distance metric, a similarity value is then produced by comparing each of the test image feature vectors.

A human perceptual experiment is used to assess the performance of the Faugeras color space in relation to the RGB and HSV color spaces. The human assessment of image color similarity provides the baseline against which each color space is compared. For a particular color space, if a high level of correlation exists between the Euclidean measurement and human evaluation, then that color space is considered a good mechanism for helping to determine image color similarity. Since the color histogram domain of content-based image retrieval attempts to abstract similarity of image colors to actual image similarity, finding a color space which corresponds closely to human perception enhances the retrieval of images with high color similarity and therefore the retrieval of similar images.

1.6 Thesis Organization

Chapter I describes the problems and benefits of utilizing content-based techniques for the retrieval of digital images. One element which helps provide accurate retrieval is the use of image color. Application of a color space based on human physiology is one possibility for trying to improve the assessment of image similarity. In Chapter II, a brief background on efficiency issues related to content-based retrieval, and the technique of color histogramming is presented. Also in the second chapter is an overview of the human visual system, and a description of the Faugeras color space. Chapter III describes the experiments performed to compare the HSV, RGB, and Faugeras color space interpretations of image similarity with those recorded by human experimentation. In Chapters IV and V, the results of the experiment are described and conclusions and recommendations are presented.

II. Background

2.1 Introduction

As discussed in Chapter I, the proliferation of digital color images has increased the need for their efficient storage and retrieval. Currently, there are many image retrieval techniques that have been either proposed or implemented. Some examples of these techniques are included in the QBIC(12), Photobook(13), and Virage image retrieval engines(14). These various systems make use of color, shape, and texture to discriminate between images. In this chapter, a subcomponent of such systems, *color histogram intersection*, is discussed in detail.

Since images are complex and not easily represented by a single word or value, their most important features must be extracted. As noted in Chapter I, the first image retrieval systems used descriptive keywords to capture the essence of an image and describe an image's content(2). These words, or *metadata*, could be used for retrieval by utilizing a string matching search. Unfortunately, using a metadata solution is feasible only when the number of images is small (hundreds or thousands of images)(2, 13). Today, image databases are expanding rapidly, and use of human workers for image annotation is wasteful. Also, with such large databases (millions of images), it is hard to describe image content by a list of words. Words are not able to capture characteristics such as texture and complex color combinations. Preservation of image information provides the ability to construct powerful search queries. Many researchers and businesses realize that while certain metadata is essential for efficient retrieval, automated indexing systems based on image content are desperately needed.

A variety of ideas have been presented that automate the process necessary to convert an image into a form for efficient retrieval (2, 12, 13, 14, 15, 16). Color-histogramming is one technique that attempts to produce simplified vectors which are representative of the original image. Each position in the vector represents one possible color from the image and the value contained in the position is a measure of how many pixels in the image are of that particular color. Image features

such as color provide a non-textual method for assessing similarity. Future sections describe the histogramming technique and also introduce improvements that can be made to the basic algorithm.

Because images are meant for human consumption, retrieval methods such as color histogramming are based on individual attributes (i.e., color) of the human visual system (HVS). A review of current content-based retrieval literature (see previous paragraph) revealed that researchers now combine multiple recognition characteristics of the HVS to improve retrieval performance. While many of the newest techniques are grounded on physiological measurements and observations of human vision, a method based on models of the human visual system was not found. The Faugeras color space (10), a model of the human color vision system, is based on human vision studies and physiological research (and incorporates multiple characteristics of both) and promises to be a useful mechanism for assessing similarity of images for image retrieval. The final two sections of this chapter provide a basic overview of the human visual system and introduce details concerning the Faugeras color space.

The discussion begins with a section on indexing and indexing structures. Correct implementation of indexing is what provides efficient retrieval and therefore requires techniques aimed at reducing the dimensionality of the search space. Otherwise, search time could grow exponentially and become infeasible for user applications. A tradeoff is always made between accurate image representation and efficient retrieval. The information presented in the next section is necessary to describe the constraints imposed by indexing. Ultimately, it is the use of indexing that determines a retrieval system's performance.

2.2 Indexing

2.2.1 Overview. The use of indexing for digital imagery has the same goals as indexing performed in relational databases (namely, speed and efficiency). However, images are much larger and are not identifiable by one unique attribute (primary key). An image must be reduced to a

set of attributes that can be employed to describe its content. The process of generating those attributes is the focus of the Color Histogramming section. The present section is concerned with how the extracted vectors are used for indexing, and why the dimensionality of the attributes must be kept to a minimum. Only short descriptions of common indexing methods will be offered since these techniques are well known. Yet, it is important to understand the utility of their application and the restrictions placed on the format of vectors used for retrieval.

2.2.2 Considerations. There are four main requirements a designer attempts to satisfy when creating an image indexing system(16). The method used should:

- be fast,
- be correct,
- incur small storage overhead, and
- be dynamic.

For the method to be fast, it must eliminate sequential scans. Like relational databases, comparison with every element in the database is not practical. When the size of the image database is comparable to a relational database and an $O(n)$ sequential scan is employed, the image database will perform slower because of the increased disk I/O and computation required to determine if two elements are similar. Also, Relational Database Management Systems (RDBMSs) search for an exact match, while image databases rarely find such a match. Instead, they look for similarity, which is not easily defined. Although most current image databases are not the size of large relational databases, their sizes are rapidly growing and, as previously stated, their comparison algorithms are usually slow. Since a sequential scan is impractical for relational systems, the added costs incurred in multimedia systems require even quicker and more efficient access methods.

The retrieval of correct results is blurred by the designer's definition of *similarity*. This definition determines what is considered a *correct* response to a query. The formal definition of

correctness in this domain is the returning of all qualifying objects without any misses. From this interpretation, false positives are acceptable, yet each design strives to minimize their presence since they have an effect on total response time.

Computer resource competition requires that space overhead be kept to a minimum. Most computer systems support more than just a data management system. If the performance a new system presents does not outweigh the increase in overhead, users will find a system that more appropriately fits their needs.

Finally, the method must be dynamic. Again, this may be dependent on the application domain, but inserts, deletes, and updates will usually need to be performed in an efficient manner. If the index is applied to a static collection of images such as a CD-ROM, this requirement will not be as important.

Since there are multiple requirements that need to be satisfied simultaneously, tradeoffs must be made between each to optimize performance. The dimensionality of vectors used for indexing (e.g. number of attributes) has the most profound effect on performance. This has been referred to as ‘the dimensionality curse’ in (16). As their size grows, the vector more accurately describes the image but increases overhead and slows look-up time. In fact, the quad tree and grid file, two common multidimensional index structures, have exponential scaleup for look-up time as dimensionality increases(17). A more efficient structure, the R-tree, is based on Minimum Bounding Rectangles(18). The R-tree and variants like the R+ and R* have been successfully tested and used for 20-30 dimension address spaces(16). An additional structure that has showed promising results was the SS-tree(19).

In the image retrieval domain, performance of indexing structures (in terms of retrieval time) is directly related to the size of the vector representing a given image. As was described, there are a number of reasons to restrict the size of this vector, the most important of these being time

constraints. These restrictions are important to keep in mind as vector extraction methods are presented.

2.3 Color Histograms

2.3.1 Overview. Color-Histograms are one way of condensing color image information into a more compact form. This compact form is usually referred to as a *feature vector*. After providing a definition for feature vectors, this section describes how the basic algorithm produces the vector and also introduces some techniques for improving efficiency and performance.

2.3.2 Feature Vectors. For efficient comparison purposes, a feature vector should be derived from an image (Figure 2.1). A feature vector is a representation of an image which carries less information, but allows very fast similarity comparisons (usually based on only one attribute - i.e., color). Figure 2.1 shows a database of images and feature vectors. As shown in this figure, the submission of a query image results in its conversion to a feature vector. This new form can easily be compared against all other vectors in the database. Any image whose feature vector is considered highly similar¹ to the query image's feature vector is returned as a query result. Describing what a feature vector is and how it is derived provides a baseline for understanding content-based retrieval.

A feature vector is simply a mathematical representation of image attributes in some n -dimensional vector space. The most common attributes used to describe an image are color, shape, texture, and relative position of objects(12, 13, 14). The vector size can increase in two ways. The most obvious way is to add more attributes. This extends the vector size by some arbitrary length which is dependent on the type of attribute (e.g., definition of a color attribute may only require a 4-byte block of memory while a texture attribute may consume 12-bytes of memory). A second case is when one attribute cannot be represented by a single number. A good example of this is spatial information (two or three dimensions - like a square/sphere). A point in three dimensional

¹The interpretation for 'highly similar' is different for each retrieval system.

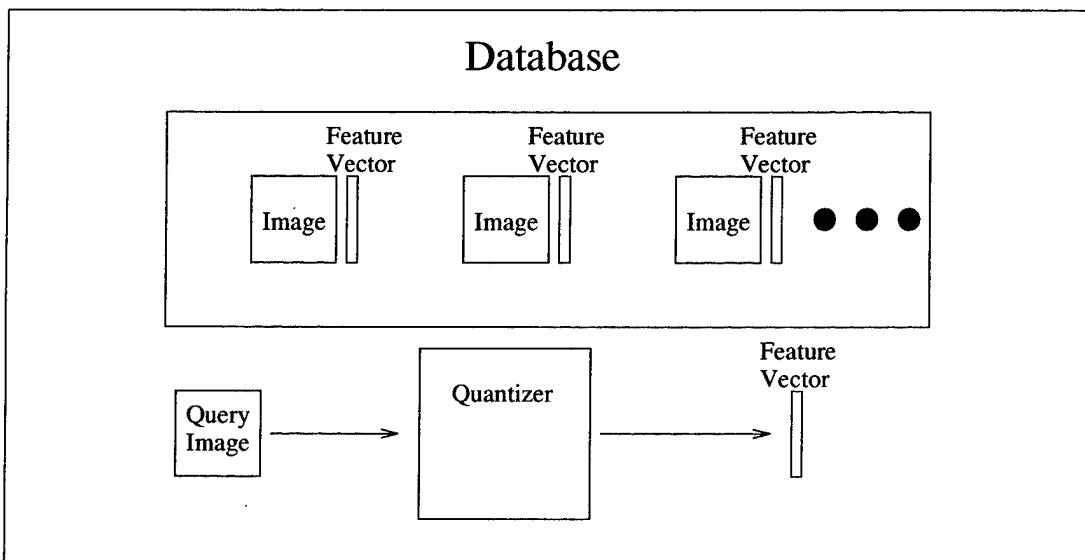


Figure 2.1 Storage of feature vectors with their corresponding image to allow quick comparisons with query images.

space can not be delineated by one number. It requires another vector whose size determines the length of each attribute.

Intuitively, the larger the vector, the more accurate the representation of the image. Unfortunately, a comparison of larger vectors usually results in slower data access. The use of domain specific comparison functions for checking similarity between images and increases in vector dimensionality result in adverse effects on search time. Figure 2.2 shows a basic model for image retrieval. There are two methods that attempt to resolve the problems inherent with content-based image retrieval. Developers can either find new indexing structures that allow efficient access to high-dimension feature vectors, or find new ways to extract a minimal amount of information that accurately describes the most important attributes of an image. Work involving indexing structures is referenced in the previous section. The rest of this section reviews color histogram-based methods for extracting minimal vectors that maximally describe an image.

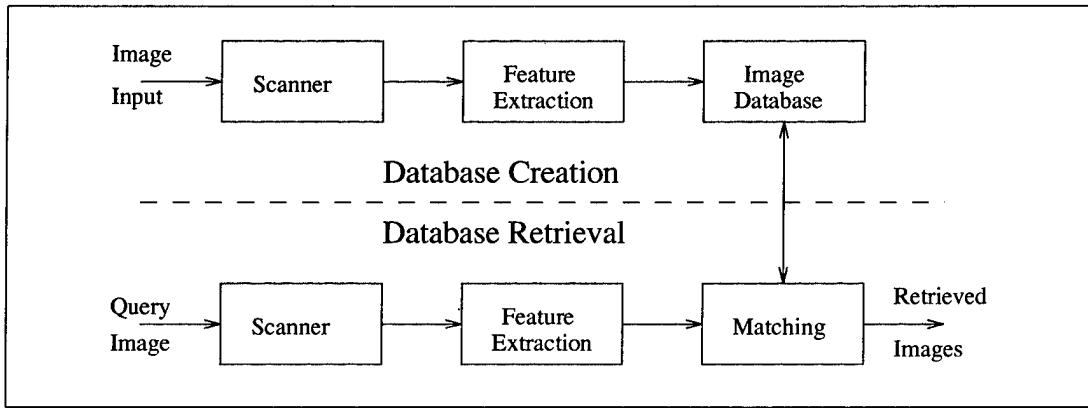


Figure 2.2 A model for content-based retrieval (2).

2.3.3 Construction and Use of a Color Histogram. Use of histograms for retrieval based on color similarities is a common strategy derived from the original work done by Ballard and Swain(9). Many improvements to this technique have been suggested (20, 21), a few of which are described in section 2.3.4. The example used here was presented in (20) and uses the RGB color space.

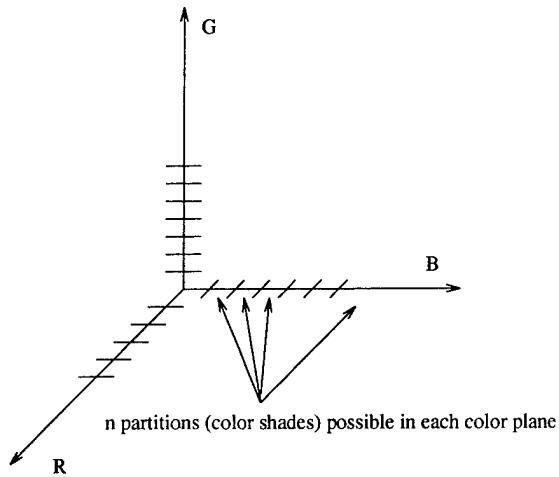


Figure 2.3 Three-Dimensional representation of the RGB color space.

The first step in constructing a color histogram is to determine the color space in which the image was encoded. With the knowledge of the three primary colors used to define the color space (in this case RGB), each axis (see Figure 2.3) of the space must be discretized (quantized) into n bins. This allows n variations of each primary color and n^3 total color combinations. Images with a 24-bit color capability result in Red, Green and Blue axes of the RGB space which each allow 256 bins. Although the high number of available colors is ideal for representing real world imagery, comparison of two such feature vectors is extremely slow. If n is picked to be a more modest value of 16, only 4096 total colors (or bins) are allowed. Fortunately, a reduced color set not only provides more efficient comparison, Ballard and Swain showed that very few color shades are necessary to maintain accurate retrieval(9). This is due to the fact that color images tend to have regions of similar colors (the green of grass, or the blue of a lake).

The histogram of the image is a vector (b_1, b_2, \dots, b_n) where each b_i holds the number of pixels from the image that correspond to that bin color. There are two ways to vectorize each image. Either a vector of length 4096 is used for retrieval (combination of all axes), or the original three vectors of length 16 (representing the color shades on each axis of the color space) are compared against each other individually. In the later case, three color histograms (one for each of the different color planes) would be produced. For example, the three histograms $(r_1, r_2, \dots, r_{16})$, $(g_1, g_2, \dots, g_{16})$,

and $(b_1, b_2, \dots, b_{16})$ could be used to represent an image in terms of the R, G, and B axes. The corresponding color axis vectors are then compared to assess similarity. The results for each axis are combined to produce an overall similarity measure. Both methods have been implemented successfully. Quantization techniques necessary to transform images with 256 bins per axis to n bins per axis (where $n < 256$) are presented in section 2.3.4. A feature vector² has now been constructed.

When retrieval is performed, a comparison is needed to determine if two images are similar. One simple metric is to compare the number of pixels in each image's corresponding histogram bins (4096 bins in the example presented above). The general similarity measure (linear) presented in (20) is:

$$d(I, H) = \sum_{l=1}^n |i_l - h_l| \quad (2.1)$$

where d is a distance function that compares the query image I to the internally stored image H . Lowercase i and h represent the value (number of pixels of that color) stored in the l th bin for the respective image. If a similar number of pixels are found in corresponding bins (for I and H), the images are considered similar. This simple distance function is useful for minimizing search time but sometimes falters on search result accuracy. Accuracy of a similarity metric is normally defined by a human's perception of the correlation between two images' color distributions. Our current inability to construct a complete model of how humans define similarity (physiologically) has resulted in the use of methods like Eq 2.1, which crudely approximate human perception, yet provide excellent performance (in terms of comparison time). The following sections describe techniques for improving accuracy while trying to maintain or reduce the required search time.

²In this case the feature vector is a color histogram

2.3.4 Methods for Improving Color Histogramming. This section contains background on the concepts that make color histogram intersection possible, and presents some examples of how retrieval accuracy can be improved. As stated in the previous section, the importance of minimizing access time requires that accuracy improvements maintain the previous level of access time performance.

2.3.4.1 Similarity Metrics. Similarity measures are a basic component of an image retrieval system and can be categorized into three groups: 1) metric measures 2) set-theoretic based measures, and 3) signal detection theory based measures. These groups can be further subdivided into measures that use crisp and fuzzy logic (22). Since algorithms in this research do not make use of the second and third categories, only crisp logic metric measures are discussed in this section. The other measures are briefly presented in (22).

Metric-based measures are frequently used to determine the color-content-based similarity between two n -dimensional feature vectors produced by color histogramming. Similarity of the images is determined by the distance between vectors. A small distance signifies high similarity while larger distances signify dissimilarity. Three measures are commonly used:

$$d_r(x, y) = \left[\sum_{i=1}^n |x_i - y_i|^r \right]^{1/r}, r \geq 1 \quad (2.2)$$

$$d_\infty(x, y) = \max |x_i - y_i| \quad (2.3)$$

where $r=1$ in equation 2.2 produces the city-block method, $r=2$ produces the Euclidean metric, and equation 2.3 is the dominance metric (22). The city-block method was used in the example of the previous section.

The Euclidean algorithm is used for image retrieval in this research. It was chosen for its mathematical tractability, and because colors in the Faugeras color space are perceptual unit distances from each other. The perceptual unit distance property of the Faugeras space allows metrics

based on distance to be used as a more accurate measure of similarity.

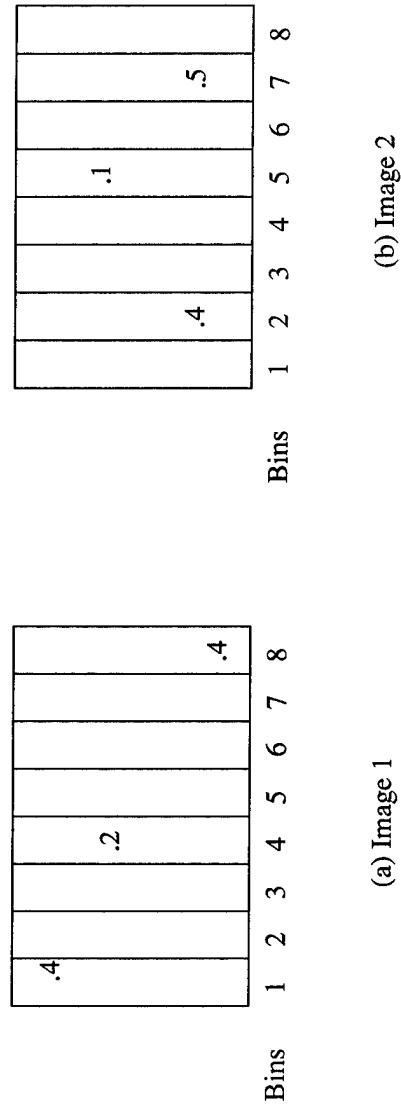


Figure 2.4 Eight bin partition of a single color space axis (3).

Although use of the Euclidean, city-block, and dominance methods may work adequately in most situations, simply comparing corresponding color bins usually limits retrieval performance since the distance metric d in equation 2.1 does not account for bins that are in close proximity to the current bin (perceptually similar colors - various shades of same color). Figure 2.4 shows

an example of two color histograms. The values in each bin represent the percentage of pixels in an image which correspond to that bin color (each bin represents one of the possible color shades defined by a color space). Since simple metrics such as equations 2.1 and 2.2 do not account for perceptually similar colors, the two images represented by the histograms are considered complete opposites. This is because the histograms are compared bin by bin. If the images do not have any color shades in common they are judged as dissimilar. Systems have solved this problem by introducing color similarity matrices. While such matrices are important for retrieval accuracy, they are not used in this research³. An introduction to color similarity matrices is given in (3).

2.3.4.2 Color Quantization Schemes. An integral component of color histogramming is color space quantization (4, 23, 24). Quantization is necessary for reducing the size of the feature vector (color histogram). In the current context, quantization is the process of dividing a color space into subsections (also called bins), thereby reducing the number of available colors an image pixel can be described by. Each bin⁴ is associated with a new output color which is an approximation of the original colors contained within that range. For example, if the Blue axis of the RGB color space was reduced from 40 bins (or colors) to 25 bins through quantization, fifteen fewer shades of blue would now exist for the definition of pixel color. The fifteen shades of blue that are lost must now be approximated by the 25 shades of blue that remain. In terms of color histograms, the vector $(b_1, b_2, \dots, b_{40})$ now becomes $(b_1, b_2, \dots, b_{25})$. After transforming the original pixel colors to the quantized pixel colors, a histogram is assembled in the same manner described in section 2.3.3.

As stated in section 2.2.2, feature vector length has a dramatic effect on access time. With color histograms, the size of the feature vector is determined by the number of bins chosen for quantization. Experiments performed in (21), which make use of known human color sensitivities,

³Accounting for adjacent bin similarity was disregarded since the inclusion or exclusion of a similarity matrix affects each color space equally.

⁴A bin defines a range of pixel intensity values from the original color space.

are necessary to determine an appropriate allocation of bin sizes for each color space axis. For example, humans are more sensitive to variations in hue, so the hue axis of the color space should be sampled more finely(21). Such tests of feature vector size provide a mechanism for optimizing the ratio of access time versus retrieval accuracy.

Improving accuracy as the length of the feature vector shrinks can be accomplished by selection of an appropriate quantization scheme for the given color space. The example presented in section 2.3.3 used a uniform quantization scheme. To quantize uniformly, the range of pixel values contained in an image (for an axis of the space) is computed and then divided by the number of desired bins. An equal range of pixel values fall within each calculated bin. Some color spaces, especially those based on color opponency⁵, do not have pixel intensity values which are distributed uniformly. Use of a uniform scheme in color opponent spaces hinders the construction of unique feature vectors. Instead, schemes like the Lloyd I (23) method are necessary because they divide the color space into a specified number of subspaces such that the resulting quantization error⁶ is minimized (21). The most common measure of quantization error is mean square error because of its mathematical tractability. The proper distribution of colors (minimizing mean square error) promotes better color matching while providing a better chance for the creation of unique feature vectors.

2.3.4.3 Color Spaces. When retrieving color imagery from a database by using color histograms, the choice of color space for similarity comparisons is extremely important(26, 27). As pointed out in section 2.3.4.2, choice of color space may determine the quantization scheme. For accurate retrieval, color representations (color spaces) must correlate well with human interpretations of color similarity. Therefore, spaces formulated from human perceptual testing make better

⁵Color Opponency is a theory based on study of the Lateral Geniculate Nucleus (LGN) that describes why humans can see red-blue shades of color but never red-green shades. The LGN, which is a structure of the brain, is thought to convert signals carried by the optic nerve into a brightness channel and two color opponent channels(25).

⁶Quantization error is the error caused by reducing the number of colors available for image display or manipulation(4).

candidates for improving color similarity-based retrieval accuracy. The following paragraphs outline the evolution of color spaces. Although standard color space schemes must be used for image display, more complicated, human perceptually based spaces are better choices when determining color similarity.

For computer users, most monitors implement the Red/Green/Blue (RGB) color system, and therefore digital imagery formats are based on the RGB color space. Yet, RGB's only connection to the end user (humans viewing the image) is that the color definition of any pixel contained in an image is a product of three values (see section 2.4.2). For 24-bit color systems, color values can range from zero to 255 and specify the amounts of red, green, and blue that are contained in each pixel (e.g. R=255, B=0, G=0 would be a definition for red). Although the system is convenient, it does not correlate well with human color perception. For example, a unit change in any of the color planes may not be perceptible by a human. This is a hindrance for determining similarity, but provides a simple solution for color display on output devices.

Fortunately, making a determination of similarity does not have to be derived from the color representation of the displayed image. Spaces based on human color perception can be and are used for similarity calculations, while the efficiency of RGB is still utilized for image display. Perceptual color spaces attempt to extract knowledge pertinent to the human definition of color similarity and eliminate any information which is extraneous to the task. As in the example above, creating a color space with unit changes that are perceptible by humans is important because a space that does not have that quality carries information unnecessary for a Euclidean similarity comparison. Since a Euclidean metric is commonly employed for comparisons, use of such spaces is an important consideration. Ideally, an appropriately designed color space only contains characteristics of color that shape a human's opinion of color similarity.

Because of a need to tune image color representation with human perception, new color spaces were researched and are now used for performing similarity calculations. One color space derived

from physically measured responses to color (by humans) is the CIE XYZ⁷(26). Although use of the CIE XYZ space is convenient, it lacks the property of perceptual uniformity - a perceptible change in color must correspond to a unit change in the values that define the space (5). This prevents the accurate use of distance metrics like those described in section 2.3.4.1.

Another color space known as the Munsell color system is a result of studies performed on humans to determine perceptual color changes(5). Munsell differs from CIE XYZ in that the measurements were gathered psychophysically (human interpretation) versus physiologically (actual physical measurements) (5). Humans were asked to give judgments of perceptual color changes and the results were used to construct the space. Yet, while distance metrics became better similarity measures, the arrangement of the space did not agree with measurements in CIE XYZ. The separation of physiological and psychophysical results for constructing color spaces finally merged with the creation of CIE LAB⁸(26). Yet, while CIE LAB is yet a better color space for testing similarity it still does not incorporate many observed (or even theorized) features of the Human Visual System (HVS). Incorporation of these new features (like the one discussed in the next paragraph) enables the definition of a more complete color space (as judged by the HVS), and should provide more accurate color image retrieval.

Currently, with advances in the knowledge of how the human visual system encodes color information, sophisticated models of human vision have been constructed (for an overview see (5)). Newer color space implementations mimic the color opponency theory and incorporate spatial as well as light sensitivities of the HVS. Color opponency researchers believe that humans convert color information into two chrominance channels (Red-Green and Blue-Yellow) and a brightness channel (see example in section 2.4.3)(25). Spatial and light sensitivities of the HVS are modeled through a combination of functions and filters. Since only minimal amounts of image retrieval

⁷CIE is an abbreviation for the Commission Internationale de L'Eclairage. They are an international color standards committee and the XYZ was a color space derived from physically measured responses of humans

⁸CIE LAB is a perceptually uniform color space designed to correlate human judgment of color similarity with a Euclidean metric's evaluation of similarity. The L, A, and B represent the three separate axes that define the color space.

research have been accomplished for these new color spaces, additional research may prove that retrieval performance (similarity based on human perceptions) can benefit from the use of color spaces which are based on more complete models of the human visual system.

2.3.5 Summary. Many improvements have been made to the original color histogramming techniques described by Ballard and Swain. Manipulation of similarity metrics, quantization schemes, and color spaces are just samples of the many proposed improvements. The lack of research with respect to color spaces generated from advanced HVS models provides an interesting avenue for further study.

Since retrieval is initiated by humans, representations of the information contained in an image may be processed better if done in a manner similar to the human visual system. In some yet unknown way, our internal representation of imagery is used to recognize the world around us. This research analyzes the affects of incorporating color HVS models into a content-based retrieval system. Color spaces derived from such models may provide better color similarity matches.

2.4 Human Visual System

2.4.1 Overview. Image storage and retrieval issues will continue to be important because of the usefulness of images to humans. Also, the utility of images as a viable information storage media is rapidly expanding as storage capacities become increasingly more cost effective. A picture is another method by which people can communicate more effectively. Amazingly, humans can recognize, interpret and understand the information contained in an image almost instantaneously. The capabilities of our visual system would be a desirable component of the ultimate image retrieval system. Unfortunately, the visual process is not completely understood and the computations and transformations made in the brain are not likely to ever be feasible on modern or future computing devices. Yet, some parts of the HVS are well understood and have been used in many application domains (e.g., image fusion(28), breast cancer detection). This section will focus on

a basic description of the HVS. This description includes examples of how scientists believe the HVS encodes, represents, and interprets images. The limitations of physical eye components, and a description of the extent of modern research will be included for completeness.

2.4.2 Image Encoding. The initial encoding phase of the human visual system involves the lens, ocular fluid and retina. This phase is where most limitations of image representation are introduced. A brief discussion illustrates the mechanics of the encoding phase.

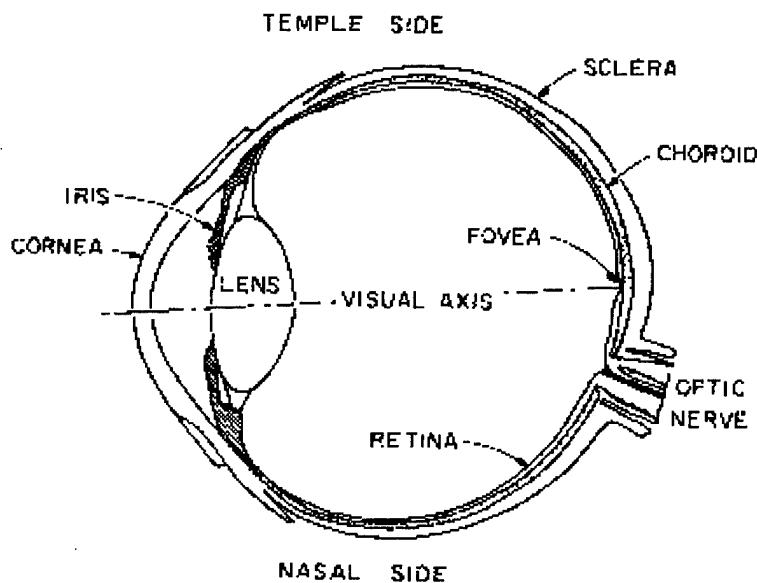


Figure 2.5 Layout of the Human Eye (4).

When light enters the eye, it is focused by the lens onto the retina (Figure 2.5). The introduction of blur by the lens is the first noted limitation of the HVS. In fact, this blur is so corrupting that Wandell(25) remarks that no person would even consider purchasing a camera with such poor optics. The human visual system sacrifices a high degree of optical accuracy in order to provide better adaptability.

After the light reaches the retina, only a very small area is capable of substantial visual acuity. This area, named the fovea, is most densely packed with light receptors. The fovea's

limited size restricts the amount of information that can be received. Therefore, only a subset of our surroundings can be examined (at a high resolution) at any instant in time.

Finally, there are only three types of receptors to encode the large bandwidth of wavelengths we consider visible light(29). Since there are only three color receptors on the retina, all colors are a combination of these three inputs⁹. Another phenomena of light absorption is the nonlinear response of the eye to increases in external light intensity(10). This phenomena occurs as the light incident on the retina is converted to a photocurrent (signal) used by the brain for image interpretation. When modeling the visual system, a nonlinear function is normally injected to approximate the nonlinear conversion that takes place between the retina and optic nerve (Figure 2.5).

Despite the limited amount of information encoded by this phase, it is enough for us to interpret the visible world. Image encoding is the most researched and best understood component of our vision. The relatively limited information provided by the encoding phase for image interpretation suggests that current computer techniques (which rely on very limited amounts of information) for image recognition or for assessing image similarity should be comparable to that of humans. Since operation of the encoding phase is well understood, researchers are now focusing on how restrictive, skewed inputs from the retina are transformed and interpreted.

2.4.3 Image Representation. As information provided by encoding proceeds through the optic nerves, it is converted in unknown (although highly theorized) ways before interpretation(25). The lack of information provided by encoding denotes a possibility that a human's internal representation of the external world greatly enhances our ability to judge attributes such as similarity. When better understood, mimicking these cortical transformations may provide more efficient and robust retrieval algorithms.

⁹The discovery of three color receptors is attributed to Young and Helmholtz and is therefore known as the Young-Helmholtz Tri-Chromatic Theory of color vision.

Transformations of the encoded signals take place in the visual pathways on their way to various parts of the brain. One focus of modern research is to understand of how these transformations affect image representation(25). Many theories have been introduced and researched, but no current model is considered to be the ‘correct’ model of the HVS. Many of the recent theories are included in human visual system models. The color opponency theory, which was introduced in an earlier section, is normally included in most modern HVS models. The theory of Color-opponency states that two chrominance channels and a brightness channel are derived from the three varieties of retinal receptor cells. This process occurs in the lateral geniculate nucleus (LGN) and results in a more efficient way to transport information from the retina to the visual cortex (25). Although opponent signals have been measured in the visual pathway, only neurons that seem to allow color transmission have been found. There is no clearly identified group of cells that transmit brightness information(25).

Multiple transformations are thought to occur in the visual pathways. The idea of color opponency is just one theory that has been presented to explain how light entering the eye is processed for interpretation by various parts of the brain. When transformations like color opponency are better understood, additional improvements in image retrieval may be realized.

2.4.4 Image Interpretation. Even after the encoded information has been transformed, the resulting retinal image is often ambiguous. Correct interpretation of the inputs is usually based on assumptions we have learned about the world around us. For example, hard objects can not pass through one another, not all types of motion are equally likely, and we live in a three-dimensional world (25). Computational methods for assessing similarity (like Euclidean distance) of images or detecting objects within a picture attempt to mimic the interpretations of the human mind. Although the least understood, the ability to define human perception will continue to be explored because of the benefits that new knowledge can provide for applications like database image retrieval.

2.5 Faugeras Color Space

2.5.1 Introduction. Previous sections of this chapter introduce the current foundations and limitations of color histogram-based image retrieval. To better understand how humans perceive and interpret images, the previous section provides a summary of the human visual system. Although computer vision research has become an integral part of image retrieval, use of models that mimic the human color visual system have not been explored for use in an information retrieval system. In this section, one such color HVS model is introduced. It is through this model that the Faugeras color space is derived. The Faugeras space provides the catalyst for the experiments outlined in Chapter 3.

2.5.2 Assumptions. Two assumptions were made when implementing the Faugeras model used in this research. First, the image must already be provided in a defined tri-stimulus space. In this case, the RGB color space was used. Second, any filtering normally performed in the HVS before the LGN color transformation stage was extracted and grouped as one filtering mechanism (contrast sensitivity function (CSF)). Justification for these assumptions can be found in (5) since the Faugeras model used in this research was originally constructed by Captain Curtis Martin as part of his dissertation.

2.5.3 Faugeras HVS Model. There are four main stages (or transformations) that the Faugeras model of human color vision accounts for(10). These stages, which are presented in the following order, include the retina color transform stage, a non-linearity stage, an LGN transform, and the CSF filters. Figure 2.6 illustrates how the model of human color vision was constructed and identifies the various stages to be discussed.

The color transform performed in the retina is represented by equation 2.4 (5).

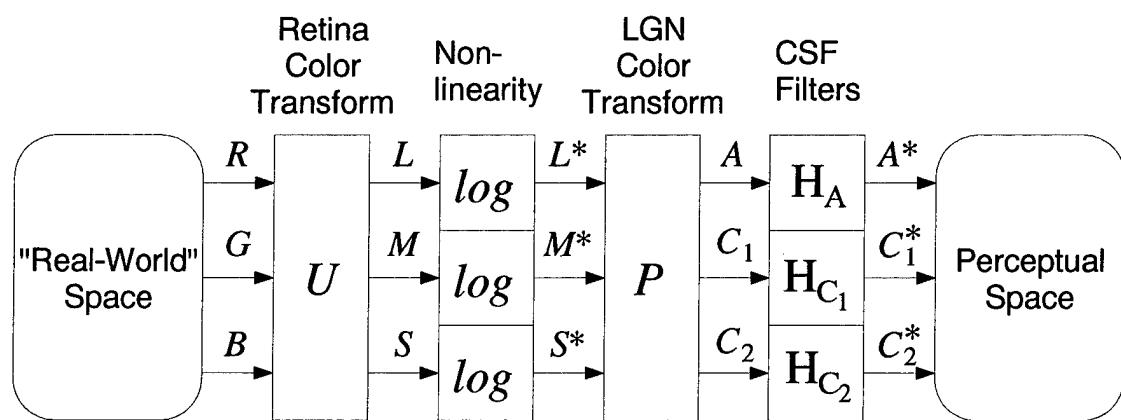


Figure 2.6 Components of the Faugeras Color HVS Model (5).

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} .3634 & .6102 & .0264 \\ .1246 & .8138 & .0616 \\ .0009 & .0602 & .9389 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}. \quad (2.4)$$

Since the input was assumed to be in terms of RGB coordinates, a linear transformation can be performed. The transform matrix of equation 2.4 represents block U in Figure 2.6. The values of U are a direct result of physical measurements of the retinal receptors' (L,M and S) reactions to red, green and blue light (5).

Next, a non-linear function is applied to the output of the retinal transform. As noted in section 2.4.2, the human eye responds nonlinearly to increases in light intensity. While not an exact match, the logarithmic function provides a reasonable approximation and is computationally simple.

The production of color opponent channels occurs in the LGN transformation. The signals from the log function (L^*, M^*, S^*) are multiplied by the matrix defined in equation 2.5.

$$\begin{bmatrix} A \\ C_1 \\ C_2 \end{bmatrix} = \begin{bmatrix} 13.8312 & 8.3394 & 0.4294 \\ 64 & -64 & 0 \\ 10 & 0 & -10 \end{bmatrix} \begin{bmatrix} L^* \\ M^* \\ S^* \end{bmatrix}. \quad (2.5)$$

Like equation 2.4, equation 2.5 was derived from psychophysically measured data. The outputs represent an achromatic channel, A, and two chromatic channels, C1 and C2. The A channel corresponds to the human perception of brightness while C1 and C2 correspond to channels containing pairs of color difference signals. In the Faugeras model, the C1 channel is composed of a Red-Green difference signal, while the C2 channel is based on the difference of a Blue and Yellow signal. The parameters of the color channels were fixed based on color change detection experiments. Colors that are just noticeably different (as decided by human subjects) are unit distances

apart. This provides a color space that conforms to an earlier definition of perceptual uniformity contained in section 2.3.4.3 (5).

Finally, a filtering stage (CSF) is applied to the A, C1 and C2 channels. These filters account for contrast and blurring effects of the HVS. In an image, contrast is a measure of differences in brightness. The application of a specific contrast function is based on studies of human sensitivity to variations in contrast. The functions used in this model are from Captain Curtis Martin's dissertation(5). They are a result of research by Mannos and Sakrison. The mathematical definition of the filters can be found in (5).

2.5.4 Limitations of Faugeras Model. There are a few qualities of the Faugeras model that cause difficulties for application in image retrieval. First, the resulting distribution of color and brightness information is Laplacian. Although this is consistent with observations of human perception, quantization of such distributions is difficult. Each image quantizes differently in order to minimize squared error¹⁰ since a non-uniform quantization is performed based on the distribution of the image being analyzed. This prevents the bin by bin comparison normally performed when implementing color histogram intersection. Bin definitions for individual images are not congruent.

A partial solution to the problem is to quantize based on a distribution of the entire database (21). An overall template can now be derived for application to each individual image. The difficulty of such a solution is that all images of the database must be present in order to construct an overall distribution. Yet, with a large enough database, any images that are added or removed will have minimal effect on the color composition of the database.

Second, even though color opponency has been well researched, the theory has not been completely validated. As stated in section 2.4.3, difference signals are transmitted from the LGN to various parts of the brain, but cells which convey brightness information have yet to be discovered.

¹⁰Again, minimization of squared quantization error corresponds to selecting the shade of color which best approximates a subset of pixel colors within an image.

2.5.5 Benefits of Using Faugeras Color Space. Although the benefits of the Faugeras color space have been noted throughout the chapter, this section provides a synopsis of the space's most important features. First, the space is perceptually uniform. Unit distances correspond to just perceptible changes in color and therefore metrics like Euclidean distance become more accurate evaluations of similarity.

Next, conversion from RGB coordinates to the Faugeras space is straightforward. This is convenient due to the commonality of RGB formats. Any RGB image can have a feature vector constructed from the Faugeras space.

Finally, the HVS model used to construct the Faugeras space accounts for human contrast sensitivities. As previously explained, contrast sensitivity functions emphasize or deemphasize differences in color (or light levels). Disposing of image features that are irrelevant to image evaluation is as important as discovering features which correspond well to human perception.

2.6 Summary

This chapter provides a background of knowledge that helps to support and explain the design of experiments and procedures carried out in Chapter III. Similarity measurements collected from different implementations of the color histogram intersection method provide a basis for comparing the performance of various color spaces to results obtained from human subjects. The Faugeras color space, which is based on human physiology, has been introduced and is analyzed in Chapter III with the histogram intersection method to determine its usefulness for assessing perceptual similarity of color images.

III. Methodology

3.1 Introduction

In this chapter, the Faugeras color space, which is discussed at the conclusion of Chapter II, is investigated as a method for improving color histogram retrieval. Retrieval improvement is based on a comparison of computer image similarity outputs to results obtained through a human perceptual test. The chapter begins with a description of the imagery, hardware, and software used for testing. Next, the methods for generating a computer-based similarity measure of two images are presented. In conclusion, the procedure for collecting human evaluations in order to rate the performance of the Faugeras color space is given.

3.2 Setup

The framework of Chapter III is based on experimentation performed in (30). In order to compare color similarity, a collection of digital color images is needed in order to replace the collection of animal forms used in (30) to test shape similarity. An appropriate data set was found on two Corel[©] CDROMs. The images contained on these CDs are a mixture of various military aircraft saved in a 24 bit Corel PCD format. From the 200 images available, 50 were chosen for their wide variations in color content since this was the feature to be used for similarity evaluation (by both the computer and humans). The 50 images were converted to a TIFF format using the UNIX *convert* function and placed in a predetermined subdirectory. To satisfy a requirement that the images be square, ImageMagick[©] (image manipulation software) was used to crop the images to their final form. When complete, the small database contained 50 128x128 pixel images stored in a 24-bit TIFF format. Although the database images were originally 256x256, a preliminary experiment for the perceptual test found that subjects preferred to judge similarity of smaller images. An image size of 128x128 was the result of that inquiry. In accordance with work done in

(30), ten test images were chosen for comparison by the computer algorithms and human subjects. The images were selected for their wide variety of hues including blues, reds, greens, and yellows.

During testing and experimentation, a Sun Sparc 20 with a 24-bit graphics card was utilized. The Matlab[©] simulation environment was used for the computer tests and presentation of the perceptual test. Any necessary image manipulation was performed with the aid of the ImageMagick software package. All relevant Matlab code for the computer tests is contained in Appendix C, and code pertaining to the perceptual experiment can be found in Appendix B.

3.3 Computer Histogram Intersection

The color-histogramming technique described in Chapter II is the basis for the algorithms constructed to evaluate the Faugeras color space. This section illustrates how the necessary color space transformations are derived, how three different types of feature vectors are generated, and how the resulting feature vectors are compared. After the composition of the color-histogramming technique has been explained, a description of how the results collected from this procedure are used for analysis is presented.

3.3.1 Color Space Transformation. For comparison purposes, three color space representations of the data are employed. In addition to the Faugeras color space, the RGB and Hue, Saturation, and Value (HSV)¹ spaces are also evaluated. The RGB color space is included because of the predominance of its use in theoretical work. The HSV space is similar to the Munsell color space described in section 2.3.4.3. The separation of hue, saturation and value distinguishes HSV from both RGB and Faugeras, and thereby provides a third distinct test space. Because of ini-

¹The HSV color space is based on the human perceptual properties of Hue, Saturation, and Value. Hue represents the basic colors of the visible light spectrum and is designated by an angular degree reading between 0 and 360 degrees (Where cyan is at 180 degrees, blue is at 240 degrees, and Magenta is at 300 degrees.) Saturation describes the vividness of a color. The value of saturation can range from 0 to 1 with values near one corresponding to complete saturation. Complete saturation of a color produces a pure spectral color, while low levels of saturation result in the perception of gray. Value represents the brightness of a color. The numerical representation of Value can also range from 0 to 1, with numbers close to 0 corresponding to darkness, while numbers close to 1 signify high levels of brightness (26).

tial test results, a fourth space was also included for comparisons. Elimination of the CSF filters described in section 2.5.3 from the Faugeras HVS model allows for the construction of the fourth color space. In the rest of the discussion, the two Faugeras spaces are referred to as the Faugeras (without CSF) and the Faugeras (with CSF) color spaces. The Faugeras (with CSF) space is derived from the original Faugeras HVS model, and the Faugeras (without CSF) space is the Faugeras HVS model without application of the CSF function.

The code used to perform the transformations from RGB to HSV and RGB to Faugeras is contained in Appendix A. Both the HSV and Faugeras color spaces are derived from an image based on RGB values. To obtain the RGB values of an image, the Matlab TIFFREAD function from the image processing toolbox is used. Once the transformations are accomplished and the data is normalized using equation 3.1, the new color space data for each image is ready to be converted to feature vectors. To normalize, μ and σ are derived from each of the color spaces. The value of μ is the mean of all data contained in the database for a particular color space plane. The standard deviation for the same collection of data is represented by σ .

$$new_datapoint = \frac{datapoint - \mu}{\sigma} \quad (3.1)$$

3.3.2 Generation of Feature Vectors. In this research, three feature vector schemes are used to compare color space performance. The schemes are based on: 1) 20 bin-per-axis uniform quantization of each color space 2) a 20 bin-per-axis non-uniform quantization of each color space based on the Lloyd I algorithm, and 3) an average of each of the color planes to produce a three-dimensional vector. Computing an average for each color plane is the simplest technique, while methods 1 and 2 produce color histograms similar to those found in (9). The non-uniform method was included to ensure that color spaces composed of non-uniform data can be accurately compared. The derivation and implementation of each scheme is discussed in the following sections.

3.3.2.1 Uniform Quantization. The first method for creating feature vectors is based on uniform quantization. As described in section 2.3.4.2, a template for building a color histogram is formed by dividing the various color spaces into n bins per axis ($n=20$ in this research). Use of quantization to produce feature vectors of length n reduces the number of representable colors, but retains a discriminating factor otherwise lost through plane averaging. This discriminating factor is the ability to compare images based on whether similar colors are present. For example, although quantization may eliminate some shades of blue, at least one shade of blue will remain. The pixel count for that one shade of blue can be compared directly with other histograms to determine if another image with a comparable number of blue pixels exists. Plane averaging loses information concerning the total number of each color shade within an image. The discriminability afforded by uniform color histograms is a major reason for their extensive use in database retrieval.

Two steps are needed to produce a uniform color histogram. First, the planes of each color space (for the whole database) are searched to find the minimum and maximum pixel intensity values. Next, these minimum and maximum values are used to define the bounds of each color plane. To quantize uniformly, the range of possible colors is divided by the number of desired colors (in this case 20). Essentially, the original color space is being quantized into a new space where each color plane now has only 20 different possible color shades (i.e. $(b_1, b_2, \dots, b_{20})$). A feature vector (color histogram) is obtained through the process described in section 2.3.3. The result is a collection of three color histograms for each image derived from the RGB, HSV, and Faugeras color spaces.

3.3.2.2 Non-Uniform Quantization. Since pixel values in the Faugeras color space are nonuniformly distributed, a nonuniform quantization scheme is required to optimize similarity comparisons. Yet, as described in section 2.5.4, application of a nonuniform quantization algorithm to individual images results in feature vectors which can not be compared for similarity. Since each image has a different distribution of pixel color, a nonuniform quantization method produces bin

ranges that are not congruent. Because of the difficulty associated with implementing a nonuniform scheme to individual images, a technique similar to the one presented in (21) is employed. Instead of quantizing based on individual images, pixel values from the whole database are used in the Lloyd I algorithm. This allows for quantization based on the color distribution of the entire database. A histogram is then constructed for each image according to this scheme.

In this research, the Matlab LLOYD I algorithm defined in the Communications Toolbox is used to optimize quantized color definitions. As input, this algorithm requires an initial *code* (codebook) and a data *training set*. The *code* is a vector whose length determines the amount of quantization. For example, if an axis with 256 color shades must be reduced to 20 color shades, then the codebook is an initial guess of what those 20 colors should be. Given the codebook, actual data points from the *training set*² are used to define a color distribution and to find the optimal partition for the quantization cells. Using the training set, the code is iteratively refined until an allowable distortion level (between original pixel color and resulting pixel color) has been attained. When complete, the LLOYD I returns an optimal partition³ and codebook vector (list of 20 optimal colors).

Quantization with this method is different from uniform quantization because an optimal approximation to the database's original pixel colors is performed. Once again, a color histogram is constructed by the method outlined in section 2.3.3.

3.3.2.3 Plane Averages. The last method for creating feature vectors is to produce an average of all pixel values for each color plane and treat the resulting three numbers (one for each plane) as a vector that represents a three dimensional point in that color space. For database retrieval, use of plane averages is the most simplistic, computationally easy comparison that can be made for color images. Unfortunately, this technique also diminishes discriminating power.

²In this research, the training set is the collection of all pixel intensity values from the entire database

³A set of 20 color ranges (containing the original 256 shades) that are to be reduced to an optimized color (one optimized color is defined for each range))

Feature vector representation of an image must allow for discriminability between images. With the averaging technique, two images with completely different colors can produce similar average values for each of the three color planes. Because a comparison of color averages may rate two images with disparate colors as similar, more false positive matches are produced (more than the uniform or non-uniform methods).

The process for generating this type of feature vector has only two steps. First, the values of all pixels for a given plane are summed and divided by the number of pixels contained in the image (Equations 3.2,3.3, and 3.4). Once this is completed for each of the three planes, the values are combined to form a vector with three positions. The example shows the creation of a feature vector for the RGB color space, but the same process is also applied to both the HSV and Faugeras color spaces.

$$R_{avg} = \frac{\sum R}{\# \text{ pixels in image}} \quad (3.2)$$

$$G_{avg} = \frac{\sum G}{\# \text{ pixels in image}} \quad (3.3)$$

$$B_{avg} = \frac{\sum B}{\# \text{ pixels in image}} \quad (3.4)$$

R_{avg}	G_{avg}	B_{avg}
-----------	-----------	-----------

Figure 3.1 Vector produced for RGB color space by the Plane Averages method

Feature vectors of the ten test images are now used in conjunction with a similarity metric to provide a third method for assessing color space performance.

3.3.3 Similarity Comparison. Assessment of image similarity is based on the distance metrics defined in section 2.3.4.1. In particular, a Euclidean metric is used to evaluate the similarity of all resulting feature vectors. Matlab code that implements a Euclidean metric can be found in Appendix C. The Euclidean distance function applied in this research does not account

for perceptual similarity between adjacent bins. As described in section 2.3.4.1, accounting for similarity between adjacent bins is important for retrieval accuracy, but was not included in this research.

The distances produced by the metric are scaled so that comparisons could be made between the various color spaces and the results obtained from the human experiment. To scale the results, the largest distance in each data set becomes the standard for images being completely dissimilar. The following method is applied to the results of the Euclidean metric to produce the final similarity output:

1. find the max distance in the data set (10x10 matrix of dissimilarity values)
2. divide all values of the data set by the value obtained by multiplying the max distance by $10/9$ - the resulting data values⁴ are between 0 and 0.9.
3. since this implementation of the Euclidean metric measures dissimilarity, the entire data set is subtracted from one to produce a measure of similarity. The change is made because the results from the human experiment will be in terms of similarity instead of dissimilarity.
4. finally, the data set is multiplied by 10 to scale the similarity values in accordance with the results obtained from the perceptual experiment.

After choosing ten test images, each is compared against all other images in the database using the method defined above. Similarity results from the distance metrics are saved in a matrix format identical to Table 3.1. For each color space, separate matrices are produced for the three quantization schemes (i.e. uniform, nonuniform, plane averages). A total of nine similarity matrices are constructed. The matrix format reflects the comparisons that are performed on the 10 test images (actual images are in table D.1). For example, the value in row 1 column 1 is the similarity value that results from comparing image 1 to itself. The similarity values can range from 1 (low similarity) to 10 (identical).

To complete an analysis of color space performance when judging similarity, data is needed from human subjects. The next section explains how the necessary data is collected.

⁴Multiplying by 10/9ths ensures that computer similarity values of 1-10 are the result of this scaling process (and not 0-10). A scale of 1 to 10 is desired because it corresponds to the similarity scale used in the human perceptual test

Table 3.1 Matrix format for storing similarity results

10	3	2	9	3	4	9	3	2	7
3	10	4	2	5	9	2	1	4	8
5	2	10	4	4	7	9	3	5	3
2	5	3	10	2	8	6	4	3	5
2	2	5	4	10	2	2	4	7	9
8	2	3	3	4	10	2	8	8	6
9	6	4	4	6	8	10	3	4	3
2	4	6	6	2	9	5	10	3	3
5	6	2	3	8	4	7	4	10	2
7	5	1	3	7	5	7	5	3	10

3.4 Perceptual Experiment

To evaluate the retrieval performance of the RGB, HSV, and Faugeras color space, a comparison mechanism is required. One way to measure retrieval accuracy is based on human evaluations of image color similarity. If an experiment is performed to collect human observations of color image similarity, the results can be used to evaluate the retrieval accuracy provided by different color spaces. The following sections describe the experiments involved in this process.

3.4.1 Color Imagery Experiment. The test subjects involved in this experiment were a collection of 34 Masters and Doctoral students and 2 Faculty members from the Air Force Institute of Technology. The task was to evaluate a subset of the 50 images contained in a test database. As described in section 3.2, ten images were chosen as the test images. Each test image was compared against a random ordering of the complete set of 10 test images (table D.1). Subjects viewed the current test image and comparison image simultaneously. A response in the range of 1 to 10 was entered via keyboard and considered a measure of the two images' perceptual color similarity. A complete description of factors each subject was made aware of prior to testing can be found in figure 3.2. These instructions were presented to each subject at the beginning of the test. When complete, subjects had compared each of the images against every other image. A total of one

hundred comparisons were performed.

GOAL - Assess perceptual similarity between digital color images

EVALUATION - The ranking of perceptual similarity will be provided with a number between 1 and 10. On this scale, the rank of 1 corresponds to a high measure of dissimilarity while a rank of 10 requires the two images to be perceived as highly similar.

CRITERIA - The determination of rank based on the above scale relies on a few important criteria:

1. The only attribute for assessing similarity is color. If the two images have similar colors (in a global sense), they are considered possible matches.
2. The images that are compared can contain completely different shapes and still be considered nearly identical (rankings close to 10). Similarity rankings should never be based on objects or shapes (for example, if both images contain an F-16). Yet two images with F-16s can be similar if their global color contents are comparable.

An analogy to keep in mind is the construction of a puzzle. One normal action when constructing a puzzle is to group pieces with similar colors. This grouping, based on the human perception of color, can be considered a filter. Similar color pieces are more likely to connect. In this experiment, you are to respond in a similar manner. Consider yourself a preprocessing filter that is deciding (based on the rank you provide) which images should be kept for further analysis.

Press any key to continue

Figure 3.2 Instructions describing the experiment.

3.4.2 Presentation of Data. For the test subjects to make similarity comparisons, the images to be compared were displayed on a computer monitor. To display multiple 24 bit images simultaneously, a Sun Sparc 20 with a 24-bit graphics card was utilized. On the monitor, a solid gray background ($R=211$, $G=211$, $B=211$) was maintained while two image windows and a Matlab window were open. Subjects were seated 18 inches from the monitor so that the image size parameter used in the Faugeras HVS was accurately represented. This parameter, which is based on distance from the monitor and image size, is derived from the following formula:

$$\theta = 2\tan^{-1}(d/D) \quad (3.5)$$

In equation 3.5, d is half the width of the displayed image, and D is the distance of the subject from the computer monitor (both in inches). The resulting angle θ corresponds to the amount of visual angle consumed by an image with respect to the subject's eye. Filters utilized in the last stage of the Faugeras color HVS model depend on visual angle for their construction. In this experiment, the visual angle was set at 5.2 degrees.

Elimination of human bias is an important consideration when performing perceptual experiments. During this experiment, several mechanisms were implemented. First, as previously described, procedural instructions were included to reduce bias and clarify the intent of the test. Second, a tutorial was presented after the instructions to anchor the range of possible similarity scores and reinforce the main points contained in the instructions. The images used in the tutorial can be found in Table D.2, Appendix D. Finally, as a result of preliminary testing on human subjects, image size was reduced from the original 256x256 pixels to 128x128 pixels. Since preliminary subjects cited a preference in evaluating the similarity of color content in smaller images, the 128x128 pixel images were used for comparison in the perceptual experiment and for evaluation by the color histogram algorithms.

3.4.3 Experimental Procedure. The experiment proceeded as follows:

1. Instructions displayed; subjects asked if they have any questions
2. Tutorial is run - three preliminary comparisons, identical to the types of comparisons made during the experiment, are used to illustrate high, low and medium color similarity between images. As described, the tutorial is meant to anchor the intent of the experiment to the 1-10 scale utilized for subject feedback.
3. Layout of experiment explained to each subject. A subject is then told how the images will be displayed, how many comparisons are being made, and when to enter their similarity response.
4. Experiment begins - Test image and first comparison image displayed (randomly chosen from collection of 10 test images)

5. While the images are being displayed, the Matlab window is idle. After the comparison image has been displayed for three seconds, the image disappears and the test subject is prompted to enter their perceived similarity measure in the Matlab window.
6. After confirmation of input, the next comparison image is displayed. This is repeated 10 times for each test image.
7. When all ten comparisons have been made, the test image is changed and the process is repeated. All ten test images are evaluated in the same manner.

As a subject cycles through all test images, their responses were recorded in vectors and saved for further processing.

3.4.4 Use of Resulting Data. Once all vectors had been collected, the results were combined statistically. Since 36 subjects provided similarity evaluations, a normal distribution was assumed⁵. For each test image, one vector was recorded per subject. Within each vector, the first value is the similarity measure between the current test image and the first test image, the second value is the similarity measure between the current test image and the second test image, and so on. There are ten test images, so each vector is of length ten. To calculate a mean value for each of the ten comparisons, the vectors resulting from the thirty-six human subjects were used. The Matlab MEAN function condensed the thirty-six vectors into a single vector which signifies the average similarity response of all subjects for a particular test image. This process was repeated for each of the ten test images. When complete, the ten resulting vectors were combined to create a 10x10 matrix similar to the matrix assembled in section 3.3.3. The mean value obtained for each similarity evaluation is used as a baseline for evaluating color space performance. Results obtained from the computer techniques described in section 3.3.2 can now be compared against the recorded similarity preferences of humans.

3.5 Summary

This chapter describes both a method for generating computer based color similarity results, and an experiment for collecting color similarity results from human subjects. The experimental

⁵usually 30 samples are enough to assume a normal distribution (31)

results were collected in order to verify the potential of the Faugeras color space for improving retrieval accuracy when using color histogram intersection. The next chapter presents the data that was collected and offers some analysis of the Faugeras color space's performance.

IV. Results

4.1 Introduction

This chapter presents the outcomes from the color histogram comparisons and human experiment described in Chapter III. First, the human perceptual experiment is analyzed to determine if any bias may have skewed the results and to provide a statistical look at the reliability of the data. The remainder of the chapter presents the results obtained for each color space. The results are subdivided into three sections based on the plane averages, uniform and non-uniform feature vector techniques described in chapter 3.

4.2 Analysis of Results from Perceptual Experiment

As described in chapter III, a mean similarity score for each image pairing was calculated. The 10x10 matrix represented by Table A.1 contains all similarity results for the ten test images. The value in position $r_{i,j}$, where r is the 10x10 matrix, signifies the similarity value that resulted from human subjects comparing image i to image j .

Because 36 subjects were sampled, the distribution of responses was assumed to be normal(31). Using this assumption, a method described in (31) is employed to determine the confidence with which the population mean μ is within a given distance of the sample mean \bar{x} . The method is based on the sample mean \bar{x} , sample standard deviation s , and sample size n of the experimental results (Equation 4.1). The purpose for using this technique is to gain statistical validation for the reliability of the data collected. The statistics show whether the mean values of the human similarity matrix would be expected to remain stable if further random samples are taken.

$$z = \frac{\bar{x} - \mu}{s/\sqrt{n}} \quad (4.1)$$

A confidence interval for the population mean is defined in equation 4.2.

$$(\bar{x} - z_{\alpha/2} \left(\frac{s}{\sqrt{n}} \right), \bar{x} + z_{\alpha/2} \left(\frac{s}{\sqrt{n}} \right)) \quad (4.2)$$

In this case, a z value was determined for each entry in the human perceptual matrix by applying equation 4.2 and setting the bounds of the confidence interval to a distance of ± 0.5 (distances relative to 1-10 scale used during experiment). A distance of 0.5 was chosen because such a change would not have a significant effect on correlation with the computer results. For this reason, this distance represents an acceptable level of change.

$$0.5 = z_{\alpha/2} \left(\frac{s}{\sqrt{n}} \right)$$

Since the sample standard deviation s and sample size n are known, z can be derived. This value of z determines the proportion of future samples (including the current sample) that will contain the true mean, μ , within a distance of ± 0.5 of the sample mean. This technique was employed for all 100 mean value entries in the human perceptual matrix. The table of resulting z values is found in Table A.2. The lowest confidence for a population mean not lying within the interval for the current and future samples was the comparison of image 7 to image 8. The computed z value of 1.40 corresponds to 83.8% confidence that this is a sample which contains the real mean μ within a distance of ± 0.5 from the sample mean, \bar{x} .

Although there is relatively high confidence that subsequent samples will produce a similar human results matrix, a few of the comparisons did exhibit high variances. Some possible reasons for these high variances include bias due to the contents of the images used, the learning curve of the human subjects and testing official, and human subject misinterpretation of the experiment instructions. The next few sections examine each of these possibilities.

4.2.1 Bias Due to Content of Images. As mentioned in Chapter III, the 10 test images were chosen based on their variations in color content. In contrast, the color of images used in the tutorial to anchor the similarity scale were dominated by several hues of blue (see table D.2).

Baselining the experiment with an image whose dominant color was blue may have introduced bias into the interpretation of the similarity scale. Table A.3 shows that the variance for comparisons between both images 1 and 2 (columns 1 and 2) and the rest of the data set are relatively low. Figure 4.1 is included to show the predominance of blue in the two test images.

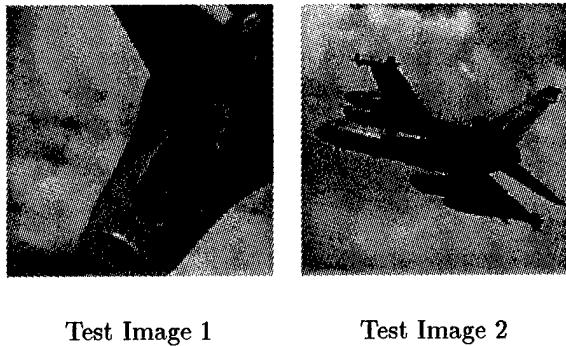


Figure 4.1 Test Images Dominated by Blue

Data present in the table of comparison variances (Table A.3) may also illustrate how the tutorial example of low similarity has affected the similarity scores. The example of dissimilarity presented in the tutorial compares an image of a plane in a bluish background to a group of planes in the early night sky (some reds and yellows which are dominated by black). This example is very similar to the comparison made between images 1 and 10 in the experiment (Figure 4.2).

As expected, Table A.3 shows low variability for the comparison between test images 1 and 10. When comparing image 10 to image 1, subjects consistently chose a similarity value of 1 or 2. In contrast, comparisons using images 3, 7, and 8 (columns 3, 7 and 8 or table A.3) all had high variances. In addition to their color contents differing from those used in the tutorial, the many shades of yellows, browns, and greens of images 3, 7, and 8 do not allow for the sensation of one overpowering global color (images found in table D.1). Without one underlying global color, the high variances suggest that similarity comparisons become much more difficult.

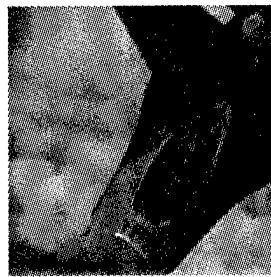
Alternately, a completely different source of bias could also have been the cause for high variances involving images 3, 7, and 8. The comparisons involving images of different brightness



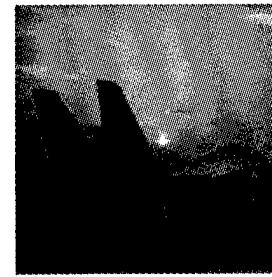
Tutorial Test Image



Example of Dissimilarity



Test Image 1



Test Image 10

Figure 4.2 Example of Possible Bias Based on Images used in Tutorial

levels (image 8 vs. images 4,5,6 and 7) suggests that there was difficulty in deciding similarities between images of varying brightness. In fact, it seems that if there is either no overpowering image color present or if no similar color shades are contained in an image, brightness levels are a secondary mechanism for assessing similarity. The results suggest that humans vary greatly in their interpretation of similarity based on brightness. Some subjects may have emphasized color content more than brightness, and others may have done just the opposite.

An additional cause of the variance (and related to the previous example) exhibited by images 3, 7, and 8 may have resulted from human interpretation of similarity between two different color hues. For example, how similar is green to blue? Such comparisons could produce a high amount of variability since the subjects were not provided with a method for deciding the similarity of different colors.

4.2.2 Bias Caused by Misinterpretation of Instructions. Another possible source of bias is instruction misinterpretation. The intent of the experiment may not have been conveyed in the instructions and tutorial. Since the experimental task was not an action that is consciously performed each day, individuals may have interpreted the instructions differently. The person administering the experiment could only explain misunderstandings presented to them by the subject. In addition, there was a learning curve involved with preparing each subject to perform the experiment. A few trials were needed before the presentation of instructions, tutorial, and an explanation of the experimental procedure were standardized.

Once the experiment had begun, each subject discovered a unique way for converting an internal similarity measure to a 1-10 evaluation. After two or three tests, people felt they could respond much quicker and had a better mental picture of what similarity ranking to assign to a pair of images. Unfortunately, since the order of image presentation was random for each subject, the actual affects could not be analyzed. An interesting idea for further study would be to determine if the variance for a comparison where image i is the test image and image j is the comparison

image differs drastically from the opposite situation (image j is the test image while image i is the comparison image).

4.2.3 Other Possible Biases. A number of other biases may have affected the results. First, although no subject complained of boredom or fatigue, 100 comparisons may have affected a subjects desire to provide unbiased answers. Second, comments were made by several subjects expressing their discouragement with the time (3 seconds) that comparison images were displayed. Since each response was a subject's initial feeling of image similarity (analysis of image color was not desired), many subjects believed that 3 seconds was too long to view the pair of images.

Next, some subjects did not maintain the 18 inch viewing distance (some leaned forward, others assumed a reclined position). Since the measurement of visual angle is based on a viewing distance of 18 inches, and the contrast sensitivity filter (CSF) depends on this value, correlation with the Faugeras results may have been biased.

Finally, as mentioned in Chapter III, image size was reduced to 128x128 pixels because of a noted preference for making similarity comparisons. Image size may still have played a role in biasing responses. Since only one image size was used, a comparison could not be made to determine if changing the size increased or decreased the variability of subject responses.

4.2.4 Summary. Because human subjects were involved, bias was a concern that the experimental setup in Chapter III attempted to minimize. In spite of these attempts, some of the results still contain high variability. Various explanations have been presented to explain deviant data. The variance in responses seemed to be affected most by the images selected to anchor the similarity scale. Further improvements can be made to future implementations of the experiment (and therefore increased confidence in the results) by incorporating changes based on comments from each of the previous sections.

4.3 Correlation Results for Computer Histogram Intersection

4.3.1 Pearson Correlation.

To measure the relationship between image color similarities observed and recorded by humans and the similarities derived from the various color histogram techniques, Pearson r values are calculated in sections 4.3.2-4.3.4. An r value describes the degree of linear relationship between two sets of data. In this case, the two sets of data were the mean values gathered from the human experiment and the individual color space similarity matrices produced from the comparison of color feature vectors. Pearson r values can range from -1 to 1 with -1 representing a perfect inverse linear relationship, +1 representing a perfect positive linear relationship, and 0 representing no linear relationship between the data sets. The Pearson Correlation function found in Microsoft's Excel was used to generate the r values contained in the current chapter.

To determine if the difference between r values is significant (i.e. the correlation shows that one color space performs better than another), a confidence for the difference between correlations was constructed. The method is derived from an example found in (32). To illustrate differences in color space performance, the process has been carried out on selected pairs of r values in each of the following sections. Each interval shows the confidence with which an r value can be considered significantly different from another. As previously stated, the sole purpose of the process is to confirm whether or not certain color spaces perform better when used as a piece of color histogram-based image retrieval.

4.3.2 Plane Average Results.

This section describes the results obtained by implementing the Plane Averages technique described in section 3.3.2.3 on each of the four color spaces. As noted in chapter 3, a 10x10 matrix of similarity comparisons was the result for each color space. Figures A.1 and A.2 in Appendix A contain the similarity matrices obtained for each of the color space representations of the images. Application of the Pearson Correlation resulted in the r values shown in Table 4.1.

Table 4.1 Pearson r Values for Plane Averages Method

	RGB	HSV	FAUGERAS WITHOUT CSF	FAUGERAS WITH CSF
Human Results	.722	.7788	.739	.736

As described in section 3.3.2.3, finding a space that improves accuracy for the plane average technique is very desirable. Unfortunately, taking a plane average will always suffer from discriminability problems since images with completely different color content can be judged as similar. The confidence levels shown in Table 4.2 suggest no significant difference in r values. In fact, there is only about 60% confidence that the HSV and RGB r values are different.

Table 4.2 Confidence that there is a Difference Between RGB and HSV r Values

Confidence Level	90%	80%	70%	60%	50%
Distance between r values	-.099	-.047	-.0112	.0158	.0398

From these results, it appears that the choice of color space does not seem to affect the retrieval accuracy of similar images. The Faugeras spaces perform on a level comparable to both HSV and RGB. As noted above, the entries in Table 4.2 resulted from considering the r values obtained for the RGB and HSV color spaces. These r values were chosen as an example because the distance between them was greatest thereby giving the best chance for a significant difference to be found.

The similarity matrices in figures A.1 and A.2 show that each color space's results overpredict the amount of similarity between images with respect to the human responses. This overprediction is caused by two factors:

1) Averaging allows images with different colors to still be similar. Since plane averaging does not maintain a count for individual colors, the similarity result may not be based on images which are composed of high numbers of the same color pixels. Therefore, a high similarity mark can be produced that does not reflect how a human observer would respond.

2) The images do not have a majority of pixels with hues that widely differ from image to image. For example, there are not many pictures with R, G and B values all close to zero. Therefore, when averaged, the vector distance between most images is small. Only images with a large difference in their pixel intensity values with respect to another image can produce averages which are great distances from other vectors. This can be seen best by observing the similarity between images 1-9 and image 10 in each color space (figures A.1-A.2). Image 10, which has a high number of deep red and orange pixels, is a large distance from the other plane average feature vectors and therefore receives a low similarity mark. This conclusion can be verified perceptually by observing the complete collection of test images in Appendix D.

As described in section 3.3.2.3, an average eliminates the possibility of comparing images based on the presence of particular colors. Further sections show that this may be a mistake when retrieval performance is based on how humans evaluate image similarity.

4.3.3 Uniform Feature Vector Results. The results for the uniform technique described in section 3.3.2.1 are shown in Figures A.3 and A.4. Once again, Pearson r values are obtained to determine the correlation between color space representation and human perception. Table 4.3 contains the r values calculated for each color space.

Table 4.3 Pearson r Values for Uniform Method

	RGB	HSV	FAUGERAS WITHOUT CSF	FAUGERAS WITH CSF
Human Results	.755	.91	.89	.725

An initial look at the r values in table 4.3 indicates superior performance for the HSV and Faugeras (without CSF)¹ color spaces. To help provide assurance that the difference between the RGB and Faugeras (without CSF) r values is significant, a confidence interval for their difference was constructed. The results in table 4.4 show 99.5% confidence that the Faugeras (without CSF)

¹CSF stands for contrast sensitivity function - refer back to section 3.3.1 for the definition of the Faugeras (with CSF) and Faugeras (without CSF) color spaces. A description of the CSF can be found in section 2.5.3

space correlates better with human results than the RGB color space. A similar analysis of the HSV and Faugeras (with CSF) was not needed since the distance between these r values was even greater than the distance used to generate the results in table 4.4.

Table 4.4 Confidence That There is a Distance Between RGB and Faugeras (without CSF) r Values (and therefore differences in performance)

Confidence Level	99.9%	99.5%	99%	95%	90%
Distance between r values	-.0301	.0398	.07213	.1605	.2049

The only recognizable difference in color representation between the test spaces is that both HSV and Faugeras are defined in terms of hue, saturation and value. The HSV space is defined directly by these concepts of human perception while the Faugeras space's axes can be used in conjunction to derive the same attributes. In contrast, as described in section 2.3.4.3, the RGB space is not derived from human color vision characteristics. In fact, the correlation value in table 4.3 supports the stance that RGB is not the best color space for judging similarity. The improvement of performance for color spaces separated into hue, value and saturation has been noted in other articles(21). The use of such spaces in this research seems to be the most plausible explanation for improvements in correlation. These results suggest further support for the use of color spaces derived from human perception when judging similarity.

The Faugeras color space (with CSF) produced the most unexpected results. The modeling of human contrast sensitivity by the CSF stage was expected to improve correlation. Instead, judgements made in this color space were no better than those made using RGB. Since the pixel values of the Faugeras (with CSF) space are distributed nonuniformly, it was thought that a uniform feature vector may have reduced performance. The unfavorable results of section 4.3.4 show that this hypothesis was false.

Another important observation is that the uniform method for feature vector production provides much better correlation with the human data than the plane averages method. Since the

uniform technique compares images color by color (albeit from a reduced color set), two images with a lot of blue will receive high marks for the similarity between them. The high correlations obtained from experimentation may suggest that humans make their similarity estimate based on a select few colors contained in each image. The wide variances (discussed in section 4.2.1) for images that lacked dominant colors also supports this stance.

4.3.4 Non-Uniform Feature Vector Results. Section 3.3.2.2 presents a method for accurate comparison of non-uniformly distributed color spaces. The Euclidean distances computed for this methodology are shown in Figures A.5 and A.6. Pearson r values (see table 4.5) are again calculated to determine the correlation between the data in Figures A.5 and A.6 and Table A.1.

Table 4.5 Pearson r Values for Non-Uniform Method

	RGB	HSV	FAUGERAS WITHOUT CSF	FAUGERAS WITH CSF
Human Results	.755	.90	.88	.68

As described in section 3.3.2.2, a non-uniform scheme was needed to optimize similarity comparisons (because of non-uniform pixel intensity distributions for the Faugeras (with CSF) color space). With a uniform scheme, most color information was clustered in the three or four bins situated around the pixel intensity value of zero in the Faugeras (with CSF) color space. A non-uniform scheme distributes the information allowing for better discriminability (and therefore better similarity comparisons).

The results of implementing this method are identical to those of the uniform scheme. Although the r values in Table 4.5 are smaller than those in Table 4.3, statistically the differences are insignificant. Since a non-uniform method should have increased correlation and didn't, the amount of overhead required to implement non-uniform quantization is not justified by the absence of performance gains in the results.

Two major components of the nonuniform technique may have produced the undesired outcome. First, only 20 bins were used for similarity testing. This may not have been enough bins to realize a performance gain. The pixel intensity distributions for the Faugeras (with CSF) color space are highly similar for each image. The small number of bins might not have provided the discriminability power necessary to produce accurate similarity evaluations.

Second, the non-uniform template was constructed based on the distribution of pixel intensity values from the database of 50 images. Although this technique was presented in (21), no performance results were offered in that paper. Applying an overall template to individual images may not be a plausible solution. The results gathered in this research support that conclusion.

4.4 Summary

This chapter presents the results obtained from comparing the human perceptual test with similarity results derived from three types of computer generated feature vectors. The feature vector results are correlated with the human test results to evaluate the retrieval accuracy provided by the Faugeras color space. The resulting correlation measures suggest that the HSV and Faugeras (without CSF) perform best (they most closely mimic how a human evaluates color image similarity). The original Faugeras space, which incorporates the use of CSF filters, performs on a level comparable to the RGB space. This is an unexpected result which is discussed further in chapter five.

V. Conclusions

Chapter V presents some conclusions based on observations noted in Chapter IV. Specifically, the performance of the Faugeras color space is discussed. The chapter closes with an overview of future research recommendations.

5.1 Performance of the Faugeras Color Space

The problem statement as presented in Chapter I is:

Does the Faugeras color space, when used as a component of color histogramming, help provide better correlation with the human perception of color image similarity than the RGB and HSV color spaces?

The next two sections describe whether better correlation was achieved, and present a number of interesting results obtained from the research.

5.1.1 Faugeras Color Space (with CSF). As explained in section 3.3.1, two versions of the Faugeras color space were used in this research. This section describes the performance of the space for which the research was originally envisioned (the Faugeras (with CSF)).

The most notable observation from the results of Chapter IV is how the use of CSF filters produced such poor correlation between human similarity measures and the similarity values produced by the uniform color histograms. This was an unexpected result since the modeling of human contrast sensitivity provides the Faugeras color space with another documented response of the human visual system. Yet, because the subjects in the human perceptual test were instructed to ignore image objects (and therefore edges), inclusion of the CSF (which emphasizes edges) was inappropriate. Therefore, color models which ignore the effects of contrast should more closely mimick the process used by subjects (in the experiment) to evaluate image similarity. The poor performance

of the Faugeras (with CSF) color space and great performance of the Faugeras (without CSF) color space supports this conclusion.

5.1.2 Faugeras Color Space (without CSF). The Faugeras (without CSF) color space was included in the research because of the high correlation values it produced during preliminary experimentation. The previous section describes why these values are much better than those produced for the Faugeras (without CSF) color space. Correlation results for this space were nearly as good as those obtained for the HSV space. As described in section 4.3.3, the only distinction between the HSV and Faugeras (without CSF) and the other color spaces was how they defined color in terms of properties such as hue, value, and saturation. Further research should be performed to determine if color spaces based on the same concepts as the HSV and Faugeras (without CSF) color spaces are best suited for evaluating color image similarity.

5.2 Recommendations for Further Work

As a result of this research, a number of areas should be explored further. First, during experimentation the length of each uniform color histogram was held constant ($n=20$). Varying the length of this vector may identify an optimal length for maximizing correlation. Knowledge of an optimal length can then be used when constructing a database retrieval system which utilizes color histogramming.

Next, other color spaces could be evaluated. The CIE-Lab color space, which is based on physiological and psychophysical research, would be a good candidate for use as a new test space. In addition, the Munsell color space should be evaluated because of its similarity to the HSV space. This would provide further proof that color spaces based on perceptual concepts like hue and saturation are best for evaluation color image similarity.

Also, the human experiment can be performed again based on the knowledge gained from the first implementation. The images used for human evaluation could be chosen based on pre-

liminary similarity results obtained from applying the uniform color histogram technique on the entire database. This would eliminate the need for test images to be chosen by the administrator of the experiment (eliminates bias introduced by a person selecting images). The suggestions for eliminating bias in sections 4.2.1-4.2.3 also need to be implemented. The application of these simple techniques may reduce the variability of human responses and therefore provide a more accurate baseline for evaluating color space performance.

Finally, although various color spaces have been compared based on their ability to provide accurate retrieval, the literature does not contain an example of performance measurement via similarity results produced by humans. Future work is needed to determine the effectiveness of the method described in this thesis for assessing color space performance with respect to database image retrieval.

5.3 Summary

The results produced in this research suggest that the Faugeras color space is a poor perceptual space for judging the similarity of images based on color. Oddly, removal of the CSF filter appears to yield large improvements in performance. Although this research provided an initial look into the performance of color spaces (for color histogram image retrieval), further work is necessary to support the resulting conclusions.

Appendix A. Similarity Matrices

Table A.1 Similarity Values Obtained from Human Perceptual Experiment.

9.97	8.38	2.64	6.39	3.46	2.99	3.43	3.49	2.17	1.33
8.38	9.97	2.85	6.67	3.22	2.89	2.97	2.85	2.11	1.61
2.64	2.85	9.97	2.68	6.44	4.07	2.83	3.47	3.82	2.72
6.39	6.67	2.68	9.97	3.60	3.65	4.42	4.53	1.99	1.67
3.46	3.22	6.44	3.60	9.94	5.36	3.90	5.25	3.31	2.35
2.99	2.89	4.07	3.65	5.36	9.92	7.11	5.25	2.04	1.60
3.43	2.97	2.83	4.42	3.90	7.11	9.72	5.28	2.15	1.38
3.49	2.85	3.47	4.53	5.25	5.25	5.28	9.97	3.75	1.94
2.17	2.11	3.82	1.99	3.31	2.04	2.15	3.75	10.00	7.28
1.33	1.61	2.72	1.67	2.35	1.60	1.38	1.94	7.28	9.83

Table A.2 z Values Obtained for Human Perceptual Results.

18.00	2.54	1.83	1.67	1.69	1.92	1.88	1.82	3.56	4.97
2.77	18.00	1.86	2.49	1.80	2.27	2.44	2.06	3.48	3.32
1.94	2.05	18.00	2.02	1.76	1.77	1.70	1.84	1.90	1.93
1.80	1.97	2.31	18.00	1.92	1.61	1.42	1.50	3.03	3.41
1.77	1.90	1.46	1.76	12.91	1.41	1.42	1.58	1.62	2.20
2.15	1.91	1.56	1.70	1.61	8.14	1.83	1.46	3.73	3.60
2.02	1.96	1.86	1.79	1.58	1.66	1.80	1.40	2.71	6.00
1.75	2.19	1.63	1.71	1.53	1.58	1.40	18.00	1.53	2.55
2.44	2.64	1.61	2.75	1.60	2.12	1.78	1.50	Inf	1.84
3.89	2.93	1.76	2.60	1.72	2.64	3.18	2.05	1.74	3.00

Table A.3 Measured Variances for Human Perceptual Results.

0.02	1.39	2.67	3.20	3.12	2.42	2.53	2.70	0.70	0.36
1.17	0.02	2.58	1.45	2.77	1.73	1.50	2.11	0.73	0.81
2.37	2.13	0.02	2.19	2.87	2.84	3.11	2.65	2.49	2.41
2.76	2.30	1.68	0.02	2.42	3.45	4.40	4.00	0.98	0.77
2.84	2.48	4.21	2.88	0.05	4.47	4.42	3.57	3.41	1.85
1.93	2.46	3.68	3.10	3.45	0.13	2.67	4.17	0.65	0.69
2.18	2.33	2.60	2.78	3.57	3.25	2.77	4.56	1.22	0.25
2.93	1.87	3.36	3.04	3.82	3.57	4.53	0.02	3.79	1.38
1.50	1.28	3.45	1.18	3.50	2.00	2.81	3.97	0.00	2.65
0.59	1.04	2.88	1.32	3.03	1.28	0.88	2.13	2.95	1.00

10.00	7.90	5.64	9.15	7.38	7.80	6.75	4.02	4.08	2.86
7.90	10.00	6.45	7.45	6.48	6.30	4.67	1.93	2.18	1.50
5.64	6.45	10.00	5.89	7.01	6.00	3.40	1.00	2.05	2.73
9.15	7.45	5.89	10.00	8.00	8.34	6.97	4.39	4.72	3.68
7.38	6.48	7.01	8.00	10.00	8.85	6.26	3.98	4.91	4.78
7.80	6.30	6.00	8.34	8.85	10.00	7.36	4.91	5.50	4.67
6.75	4.67	3.40	6.97	6.26	7.36	10.00	7.13	6.58	4.17
4.02	1.93	1.00	4.39	3.98	4.91	7.13	10.00	7.77	4.49
4.08	2.18	2.05	4.72	4.91	5.50	6.58	7.77	10.00	6.71
2.86	1.50	2.73	3.68	4.78	4.67	4.17	4.49	6.71	10.00

RGB

10.00	8.85	5.12	9.01	6.51	6.88	7.93	5.63	4.81	2.04
8.85	10.00	5.18	8.33	6.19	6.45	6.95	4.78	4.18	1.81
5.12	5.18	10.00	5.53	8.15	7.67	5.68	5.27	6.28	2.77
9.01	8.33	5.53	10.00	6.96	7.37	8.08	6.39	5.02	1.58
6.51	6.19	8.15	6.96	10.00	9.49	7.48	6.78	7.28	3.06
6.88	6.45	7.67	7.37	9.49	10.00	7.94	7.14	7.22	2.91
7.93	6.95	5.68	8.08	7.48	7.94	10.00	7.14	6.61	2.89
5.63	4.78	5.27	6.39	6.78	7.14	7.14	10.00	6.20	1.00
4.81	4.18	6.28	5.02	7.28	7.22	6.61	6.20	10.00	4.47
2.04	1.81	2.77	1.58	3.06	2.91	2.89	1.00	4.47	10.00

HSV

Figure A.1 Similarity Matrices for Plane Average Feature Vectors Produced from the RGB and HSV Spaces.

10.00	9.56	6.58	9.15	7.51	7.70	8.65	7.47	4.74	1.00
9.56	10.00	6.76	8.93	7.54	7.70	8.36	7.18	4.68	1.02
6.58	6.76	10.00	6.91	8.67	8.37	6.81	6.69	6.64	3.75
9.15	8.93	6.91	10.00	7.99	8.05	8.78	8.18	5.51	1.71
7.51	7.54	8.67	7.99	10.00	9.45	8.02	7.99	6.82	3.31
7.70	7.70	8.37	8.05	9.45	10.00	8.37	7.99	6.41	2.89
8.65	8.36	6.81	8.78	8.02	8.37	10.00	8.32	5.36	1.58
7.47	7.18	6.69	8.18	7.99	7.99	8.32	10.00	6.65	2.73
4.74	4.68	6.64	5.51	6.82	6.41	5.36	6.65	10.00	6.07
1.00	1.02	3.75	1.71	3.31	2.89	1.58	2.73	6.07	10.00

Faugeras without CSF

10.00	9.53	6.63	9.12	7.60	7.83	8.69	7.44	4.72	1.00
9.53	10.00	6.84	8.88	7.64	7.83	8.36	7.13	4.66	1.03
6.63	6.84	10.00	6.94	8.62	8.29	6.71	6.55	6.48	3.64
9.12	8.88	6.94	10.00	8.08	8.16	8.77	8.18	5.52	1.74
7.60	7.64	8.62	8.08	10.00	9.42	7.96	7.89	6.71	3.22
7.83	7.83	8.29	8.16	9.42	10.00	8.36	7.90	6.26	2.76
8.69	8.36	6.71	8.77	7.96	8.36	10.00	8.23	5.24	1.46
7.44	7.13	6.55	8.18	7.89	7.90	8.23	10.00	6.59	2.68
4.72	4.66	6.48	5.52	6.71	6.26	5.24	6.59	10.00	6.07
1.00	1.03	3.64	1.74	3.22	2.76	1.46	2.68	6.07	10.00

Faugeras with CSF

Figure A.2 Similarity Matrices for Plane Average Feature Vectors Produced from the RGB and HSV Spaces.

10.00	5.37	4.39	4.21	5.30	5.28	4.49	3.60	3.95	2.22
5.37	10.00	3.22	4.05	2.83	3.27	1.00	1.11	2.26	1.42
4.39	3.22	10.00	3.46	3.57	4.45	2.44	1.49	2.14	1.33
4.21	4.05	3.46	10.00	5.37	5.20	3.85	3.69	5.03	4.79
5.30	2.83	3.57	5.37	10.00	6.48	5.22	4.31	5.24	4.08
5.28	3.27	4.45	5.20	6.48	10.00	6.05	3.64	5.11	3.46
4.49	1.00	2.44	3.85	5.22	6.05	10.00	4.60	5.23	3.08
3.60	1.11	1.49	3.69	4.31	3.64	4.60	10.00	6.11	3.60
3.95	2.26	2.14	5.03	5.24	5.11	5.23	6.11	10.00	4.70
2.22	1.42	1.33	4.79	4.08	3.46	3.08	3.60	4.70	10.00

RGB

10.00	7.65	3.08	5.79	4.24	4.51	5.30	2.70	3.39	1.11
7.65	10.00	3.82	6.16	3.72	3.59	4.39	2.47	2.76	1.00
3.08	3.82	10.00	3.23	5.65	4.98	3.79	4.16	3.30	1.27
5.79	6.16	3.23	10.00	4.68	4.58	5.13	3.36	4.88	3.14
4.24	3.72	5.65	4.68	10.00	6.07	6.35	5.59	4.79	2.69
4.51	3.59	4.98	4.58	6.07	10.00	6.27	4.58	4.75	2.19
5.30	4.39	3.79	5.13	6.35	6.27	10.00	5.26	4.89	2.11
2.70	2.47	4.16	3.36	5.59	4.58	5.26	10.00	3.98	1.50
3.39	2.76	3.30	4.88	4.79	4.75	4.89	3.98	10.00	5.84
1.11	1.00	1.27	3.14	2.69	2.19	2.11	1.50	5.84	10.00

HSV

Figure A.3 Similarity Matrices for Uniform Feature Vectors Produced from the RGB and HSV Spaces.

10.00	7.58	1.07	5.76	3.38	3.50	4.55	2.63	2.71	2.34
7.58	10.00	1.75	6.20	3.02	3.63	3.86	2.13	2.02	2.42
1.07	1.75	10.00	3.06	4.85	5.61	2.01	3.89	3.63	1.22
5.76	6.20	3.06	10.00	6.21	5.22	5.62	4.22	4.22	2.18
3.38	3.02	4.85	6.21	10.00	6.86	5.05	4.91	5.67	2.13
3.50	3.63	5.61	5.22	6.86	10.00	5.59	4.92	4.29	1.89
4.55	3.86	2.01	5.62	5.05	5.59	10.00	4.55	4.54	1.00
2.63	2.13	3.89	4.22	4.91	4.92	4.55	10.00	4.83	1.21
2.71	2.02	3.63	4.22	5.67	4.29	4.54	4.83	10.00	4.29
2.34	2.42	1.22	2.18	2.13	1.89	1.00	1.21	4.29	10.00

Faugeras without CSF

10.00	8.90	6.14	6.73	3.04	2.91	5.20	5.61	4.43	1.00
8.90	10.00	5.90	6.78	3.20	3.03	4.87	5.22	4.95	1.10
6.14	5.90	10.00	5.00	4.51	3.21	6.62	7.26	3.78	2.17
6.73	6.78	5.00	10.00	6.21	5.68	5.08	6.86	5.29	3.65
3.04	3.20	4.51	6.21	10.00	7.60	6.89	6.15	5.74	6.52
2.91	3.03	3.21	5.68	7.60	10.00	5.92	4.44	6.37	6.29
5.20	4.87	6.62	5.08	6.89	5.92	10.00	5.94	6.23	4.68
5.61	5.22	7.26	6.86	6.15	4.44	5.94	10.00	4.37	3.94
4.43	4.95	3.78	5.29	5.74	6.37	6.23	4.37	10.00	4.76
1.00	1.10	2.17	3.65	6.52	6.29	4.68	3.94	4.76	10.00

Faugeras with CSF

Figure A.4 Similarity Matrices for Uniform Feature Vectors Produced from the Faugeras Spaces.

10.00	5.47	4.50	4.29	5.39	5.36	4.47	3.56	3.94	2.17
5.47	10.00	3.34	4.06	2.83	3.29	1.00	1.02	2.22	1.42
4.50	3.34	10.00	3.57	3.64	4.54	2.47	1.49	2.12	1.33
4.29	4.06	3.57	10.00	5.31	5.26	3.87	3.65	4.98	4.91
5.39	2.83	3.64	5.31	10.00	6.46	5.22	4.19	5.17	3.70
5.36	3.29	4.54	5.26	6.46	10.00	6.06	3.54	5.04	3.38
4.47	1.00	2.47	3.87	5.22	6.06	10.00	4.53	5.24	3.01
3.56	1.02	1.49	3.65	4.19	3.54	4.53	10.00	6.01	3.39
3.94	2.22	2.12	4.98	5.17	5.04	5.24	6.01	10.00	4.44
2.17	1.42	1.33	4.91	3.70	3.38	3.01	3.39	4.44	10.00

RGB

10.00	7.43	3.12	5.48	4.33	4.43	5.31	2.38	3.24	1.00
7.43	10.00	3.52	5.75	3.66	3.42	4.24	2.26	2.62	1.13
3.12	3.52	10.00	2.97	5.40	4.74	3.75	3.92	2.95	1.23
5.48	5.75	2.97	10.00	4.66	4.54	5.04	3.00	4.82	2.77
4.33	3.66	5.40	4.66	10.00	6.01	6.28	5.65	4.55	2.41
4.43	3.42	4.74	4.54	6.01	10.00	6.22	4.42	4.55	1.82
5.31	4.24	3.75	5.04	6.28	6.22	10.00	5.06	4.71	1.96
2.38	2.26	3.92	3.00	5.65	4.42	5.06	10.00	3.62	1.16
3.24	2.62	2.95	4.82	4.55	4.55	4.71	3.62	10.00	5.46
1.00	1.13	1.23	2.77	2.41	1.82	1.96	1.16	5.46	10.00

HSV

Figure A.5 Similarity Matrices for Non-Uniform Feature Vectors Produced from the RGB and HSV Spaces.

10.00	7.35	1.22	5.62	3.53	3.31	4.16	3.23	3.28	2.70
7.35	10.00	1.00	5.82	3.14	3.27	3.01	2.91	2.53	3.11
1.22	1.00	10.00	2.99	4.55	4.73	1.76	2.37	4.12	1.48
5.62	5.82	2.99	10.00	5.96	4.75	4.99	4.30	4.57	2.80
3.53	3.14	4.55	5.96	10.00	6.35	5.25	4.51	5.53	2.54
3.31	3.27	4.73	4.75	6.35	10.00	5.72	4.40	4.53	2.48
4.16	3.01	1.76	4.99	5.25	5.72	10.00	4.24	4.64	1.63
3.23	2.91	2.37	4.30	4.51	4.40	4.24	10.00	4.57	1.53
3.28	2.53	4.12	4.57	5.53	4.53	4.64	4.57	10.00	4.30
2.70	3.11	1.48	2.80	2.54	2.48	1.63	1.53	4.30	10.00

Faugeras without CSF

10.00	7.15	6.26	5.94	4.66	4.50	4.34	3.63	3.01	3.14
7.15	10.00	6.65	5.62	4.83	5.43	3.38	1.94	2.23	3.25
6.26	6.65	10.00	4.18	4.32	4.76	2.88	1.37	1.00	2.55
5.94	5.62	4.18	10.00	6.41	5.30	4.81	3.31	4.49	5.42
4.66	4.83	4.32	6.41	10.00	7.53	5.88	2.64	4.62	5.79
4.50	5.43	4.76	5.30	7.53	10.00	5.30	1.34	3.81	5.26
4.34	3.38	2.88	4.81	5.88	5.30	10.00	4.35	6.11	3.26
3.63	1.94	1.37	3.31	2.64	1.34	4.35	10.00	5.16	1.74
3.01	2.23	1.00	4.49	4.62	3.81	6.11	5.16	10.00	3.74
3.14	3.25	2.55	5.42	5.79	5.26	3.26	1.74	3.74	10.00

Faugeras with CSF

Figure A.6 Similarity Matrices for Non-Uniform Feature Vectors Produced from the Faugeras Spaces.

Appendix B. Experiment Matlab M-Files

EXPERIMENT

```
function experiment(subjectnum)
%
%
%
instructions;
tutorial;
perceptMatrix(subjectnum);
```

TUTORIAL

```
function tutorial
%
%
%
%
!display -title A -geometry +420+80 IMAGES/HALFI/image5.tif \&;
[testimagepid] = getpid(0,0);

!display -title B -geometry +555+80 IMAGES/HALFI/image43.tif \&;
[pid] = getpid(1,testimagepid);
fprintf('Press space bar to continue\n\n\n');
pause;
eval(['!kill ' int2str(pid)]);

!display -title C -geometry +555+80 IMAGES/HALFI/image41.tif \&;
[pid] = getpid(1,testimagepid);
fprintf('Press space bar to continue\n\n\n');
pause;
eval(['!kill ' int2str(pid)]);

!display -title D -geometry +555+80 IMAGES/HALFI/image31.tif \&;
```

GETPID

```
%
%
%
%
[pid] = getpid(1,testimagepid);
```

```

fprintf('Press space bar to continue\\n\\n\\n');
pause;
eval(['!kill ' int2str(pid)]);

eval(['!kill ' int2str(testimagepid)]);

function [pid] = getpid(switch,origpid)

if switch == 0

!ps -a > output;
!grep display output > processes;
fid = fopen('processes','r');
pid = fscanf(fid,'%d');
fclose(fid);

else

!ps -a > output;
!grep display output > processes;
eval(['!grep -v ' int2str(origpid) ' processes > finallist']);
fid = fopen('finallist','r');
pid = fscanf(fid,'%d');
fclose(fid);

end

```

PERCEPMATRIX

```

function perceptMatrix(testsubject)
%
%
%
%
%
fprintf('\\n\\n      This is the start of Test #1 \\n\\n');
fprintf('      You will have 3 seconds to make a comparison \\n');
fprintf('      before the image on the right is removed and a\\n');
fprintf('      similarity score must be entered. \\n\\n');
fprintf('Press any key to continue\\n\\n');
pause;

flag=0;
CONTROL = 1;
testimageflag=0;
x=clock;
randnum = ceil(x(6));
randperm(randnum);

```

```

testimage = randperm(10);
for k = 1:10

outvector = zeros(10,10);

filename = ['IMAGES/EXP\_IMAGES/image' int2str(testimage(k)) '.tif'];
eval(['!display -title Test\_' int2str(k) ' -geometry +420+80 ' filename ' \&']);
[testimagepid] = getpid(0,0);
compimage = randperm(10);

for i = 1:10

%must randomize input to int2str

filename = ['IMAGES/EXP\_IMAGES/image' int2str(compimage(i)) '.tif'];
eval(['!display -title ' int2str(i) ' -geometry +555+80 ' filename ' \&']);
[pid] = getpid(1,testimagepid);
pause(3);
eval(['!kill ' int2str(pid)]);

while (flag==0)
    simscore = input('Enter a similarity score from 1-10: ');
    fprintf('\n\n');

%check validity of inputted score
if (simscore == 1) | (simscore ==2) | (simscore ==3) | (simscore ==4) | (simscore ==5) | (simscore ==6)
    outvector(testimage(k),compimage(i))=simscore;
    flag=1;
else
    fprintf('Similarity score must be between 1 and 10, please enter score again.\n\n');
end
end
flag=0;
%fprintf('Press space bar to continue\n\n\n');
%pause;

end

testimageflag = 1;
eval(['!kill ' int2str(testimagepid)]);

if k~=10
    output = ['\n\n' This is the start of Test ' int2str(k+1) ' \n\n'];
    fprintf(output);
    fprintf('Press any key to continue');
    pause
end

```

```

filename = ['''TESTIMAGE\_` int2str(testimage(k)) '/subject\` int2str(testsubject) '\_results'''
eval(['save ' filename ' outvector']);

clear outvector

end

SCRAMBLE128
%
%
%
%
for j = 1:10

infile = ['IMAGES/EXP\_IMAGES/image' int2str(j)];
[R,G,B] = tiffread(infile);

x = randperm(16384);

for i = 1:16384

R(x(i))=R(i);
G(x(i))=G(i);
B(x(i))=B(i);

end

outfile = ['IMAGES/SCRAMBLED\_EXP/image' int2str(j) '.tif'];
tiffwrite(R,G,B,outfile);

end

```

Appendix C. Color Histogram Matlab M-Files
COLOR SPACE TRANSFORMATIONS

RGB2AC1C2

```
function [A,C1,C2,R,G,B] = rgb2ac1c2(image,imgsize)
[R,G,B] = tifffread(image);
[R,G,B] = replacezero(R,G,B);
[A,C1,C2] = colorhvs(R,G,B,imgsize);
```

COLORHVS - FILTERED

```
function [A1,B1,C1] = colorhvs(r1,g1,b1,imgsize)
%
% r1,g1,b1 = input (reference) image red, green, and blue planes
% r2,g2,b2 = input (distorted) image red, green, and blue planes
% imgsize = size of image in degrees of visual angle
%
% pA,pC1,pC2 = output visible difference maps for the A, C1, and C2 planes
%
% Author: Curtis E. Martin
% Date: 17 Sep 96

% Should check sizes, etc....
[N,M]=size(r1);

% Initialize variables

sA = zeros(N*M,1);
sB = zeros(N*M,1);
sC = zeros(N*M,1);

% Set parameters:
$W = 6; % Daly, p. 196
Q = 0.7; % Daly, p. 196
k1 = W^(-Q/(1-Q)); % Daly's equation 14.30
k2 = W^(1/(1-Q)); % Daly's equation 14.30
b = 4; % Daly, p. 197
s = .8; % Daly's varied from 0.7 to 1.0
beta = 3.4; % based on plot, Daly, p. 198
$

% Get HVS filters
Ha = gethvs(N, imgsize, 1);
Hc1 = gethvs(N, imgsize, 2);
Hc2 = gethvs(N, imgsize, 4);
```

```

% Perform colorimetric transformation to Faugeras color space (no filtering)
[A1,B1,C1] = rgb2faug(r1,g1,b1);

% Compute Fourier transforms
$A1 = fft2(A1)/(N^2);
B1 = fft2(B1)/(N^2);
C1 = fft2(C1)/(N^2);$

% Now apply the CSF filters
A1 = Ha .* A1;
B1 = Hc1 .* B1;
C1 = Hc2 .* C1;

% Convert back to time domain
$A1 = real(ifft2(A1))*(N^2);
B1 = real(ifft2(B1))*(N^2);
C1 = real(ifft2(C1))*(N^2);$

```

GETHVS

```

function H = gethvs(N,isize,plane)
% gethvs: build isotropic Contrast Sensitivity Function filter for application
%           in the Fourier domain.
%
% H = gethvs(N,isize,plane)
%
% N      = number of pixels in images
% isize = size of image on screen, in degrees
% plane = color plane (1=A, 2=C1, and 4=C2)
%
% H      = bandpass modulation transfer function

% The CSF filter that is returned is obtained by evaluating the functional
% form given by Mannos & Sakrison, IEEE Transactions on Information Theory,
% vol. IT-20, no. 4, pp. 525-536. It is calibrated to peak at about
% 8 cycles/deg for the A component, 4 cycles/deg for the C1 component, and 2
% cycles/deg for the C2 component.

fs = 1 / isize;
ii = ones(N,1) * [-floor(N/2):floor((N-1)/2)];
jj = [-floor(N/2):floor((N-1)/2)]' * ones(1,N);
w = plane * 0.018 * fftshift(2*pi*fs*sqrt($ii.^2 + jj.^2$));
H = 2.6 * (0.0192 + w) .* exp(-($w.^1.1$));

```

COLORHVS2 - NONFILTERED

```

function [A1,B1,C1] = colorhvs2(r1,g1,b1)
%
% r1,g1,b1 = input (reference) image red, green, and blue planes
%
% Author: Curtis E. Martin
% Date: 17 Sep 96
% Date: 20 Sep 97 - Modified by Chad A. Vander Meer
% - Removed filtering

% Should check sizes, etc....
[N,M]=size(r1);

% Initialize variables

sA = zeros(N*M,1);
sB = zeros(N*M,1);
sC = zeros(N*M,1);

% Set parameters:
W = 6; % Daly, p. 196
Q = 0.7; % Daly, p. 196
k1 = $W^(-Q/(1-Q))$; % Daly's equation 14.30
k2 = $W^(1/(1-Q))$; % Daly's equation 14.30
b = 4; % Daly, p. 197
s = .8; % Daly's varied from 0.7 to 1.0
beta = 3.4; % based on plot, Daly, p. 198

%
% Perform colorimetric transformation to Faugeras color space (no filtering)
[A1,B1,C1] = rgb2faug(r1,g1,b1);

```

RGB2FAUG

```

function [a, c1, c2] = rgb2faug(r,g,b)
%RGB2FAUG Transform an RGB image or colormap into Faugeras' uniform color
% space. [A, C1, C2] = RGB2FAUG(R,G,B) returns the three components of
% Faugeras' color space equivalent to the RGB image in the matrices R,
% G, B.
% NEWMAP = RGB2FAUG(MAP) transforms the the M-by-3 colormap MAP into
% the equivalent Faugeras colormap MAP.
%
% See also FAUG2RGB

%
% (((The matrices R, G, B, contain intensities in the range
% from 0 to 1 of the red, green, and blue components of the image. The
% intensity 0 corresponds to black, while the intensity 1
% corresponds to full intensity. ???)))

```

```

% Curtis E. Martin 29 Aug 95

if nargin == 2 | nargin > 3,
    error('Wrong number of output arguments.');
end

if nargin==1,
    rgb = r';
    a = zeros(size(r));

elseif nargin==2,
    error('Wrong number of input arguments.');
else
    if (any(size(r)==size(g)) | any(size(r)==size(b))),
        error('R, G, and B must all be the same size.')
    end
    rgb = [r(:); g(:); b(:)];
    a = zeros(size(r));
    c1 = zeros(size(r));
    c2 = zeros(size(r));
end

% Define U and P here
% U1:
%U = [.2457 .6840 .0703; .1101 .7625 .1273; .0132 .0842 .9026];
% U2:
% U = [.3634 .6102 .0264; .1246 .8138 .0616; .0009 .0602 .9389];
P = [13.8312 8.3394 0.4294; 64 -64 0; 10 0 -10];
temp = P * log(U * rgb);

if nargout == 1,
    a = temp';
else
    a(:) = temp(1,:);
    c1(:) = temp(2,:);
    c2(:) = temp(3,:);
end

```

RGB2HSV2

```

function [h,s,v] = rgb2 hsv2(r,g,b)
%Given: rgb each in [0,1].
%Desired: h in [0,360] and s in [0,1], except if s=0, then h=UNDEFINED.
h = zeros(size(r,1));
s = zeros(size(g,1));
v = zeros(size(b,1));

$n = size(r,1)^2;

```

```

for i = 1:n
    max1 = max([r(i) g(i) b(i)]);
    min1 = min([r(i) g(i) b(i)]);
    v(i) = max1;                                %This is the lightness
    if max1 ~= 0
        s(i) = (max1-min1)/max1; %Next calculate saturation
    else
        s(i) = 0;
    end

    if s(i) == 0
        h(i) = eps;
    else delta = max1-min1;

        if r(i) == max1
            h(i) = (g(i)-b(i))/delta;           %Resulting color is between yellow and magenta}
        elseif g(i) == max1
            h(i) = 2 + (b(i)-r(i))/delta;      %Resulting color is between cyan and yellow}
        elseif b(i) == max1
            h(i) = 4 + (r(i)-g(i))/delta;      %Resulting color is between magenta and cyan}
        end                                     %Convert to degrees
        h(i) = h(i)*60;
        if h(i) < 0.0
            h(i) = h(i) + 360;                %Make degrees be nonnegative}
        end
    end
end
*****

```

DB_DIST

```

Avect = [];
C1vect = [];
C2vect = [];

Afvect = [];
C1fvect = [];
C2fvect = [];

Rvect = [];
Gvect = [];
Bvect = [];

Hvect = [];
Svect = [];
Vvect = [];

for i = 1:50

```

```

infile = ['NORMALIZED_DATA/ALL_IMAGES2/image' int2str(i)];
eval(['load ' infile]);

Avect = [Avect; A(:)];
C1vect = [C1vect; C1(:)];
C2vect = [C2vect; C2(:)];

Afvect = [Afvect; Af(:)];
C1fvect = [C1fvect; C1f(:)];
C2fvect = [C2fvect; C2f(:)];

Rvect = [Rvect; R(:)];
Gvect = [Gvect; G(:)];
Bvect = [Bvect; B(:)];

Hvect = [Hvect; H(:)];
Svect = [Svect; S(:)];
Vvect = [Vvect; V(:)];

end

Amin = min(Avect);
Amax = max(Avect);
C1min = min(C1vect);
C1max = max(C1vect);
C2min = min(C2vect);
C2max = max(C2vect);

Afmin = min(Afvect);
Afmax = max(Afvect);
C1fmin = min(C1fvect);
C1fmax = max(C1fvect);
C2fmin = min(C2fvect);
C2fmax = max(C2fvect);

Rmin = min(Rvect);
Rmax = max(Rvect);
Gmin = min(Gvect);
Gmax = max(Gvect);
Bmin = min(Bvect);
Bmax = max(Bvect);

Hmin = min(Hvect);
Hmax = max(Hvect);
Smin = min(Svect);
Smax = max(Svect);
Vmin = min(Vvect);
Vmax = max(Vvect);

```

```

save 'UNIFORM_HISTOGRAMS/db_min_max2' Amin Amax C1min C1max C2min C2max Afmin Afmax C1fmin C1fmax

LLOYDTESTAUF

Afvect = [];

load 'UNIFORM_HISTOGRAMS/db_min_max2';
step = (Afmax-Afmin)/19;
code = Afmin:step:Afmax;

for i = 1:50
    filename = ['NORMALIZED_DATA/ALL_IMAGES2/image' int2str(i)];
    eval(['load ' filename]);
    NewAfvect = Af(1:16384);
    Afvect = [Afvect NewAfvect];
end

[Partition,Codebook,Distortion] = lloyds(Afvect,code,.001);
save Afoptimization Partition Codebook Distortion

```

PLANEFV1_N

```

variables = [' Aavg C1avg C2avg Afavg C1favg C2favg Ravg Gavg Bavg Havg Savg Vavg'];

for j = 1:10

    filename = ['NORMALIZED_DATA/EXP_IMAGES_2/image' int2str(j)];
    eval(['load ' filename]);

    numpixels = (size(A,1))^2;

    Aavg = (sum(A(:))/numpixels);
    C1avg = (sum(C1(:))/numpixels);
    C2avg = (sum(C2(:))/numpixels);

    Afavg = (sum(Af(:))/numpixels);
    C1favg = (sum(C1f(:))/numpixels);
    C2favg = (sum(C2f(:))/numpixels);

    Ravg = sum(R(:))/numpixels;
    Gavg = sum(G(:))/numpixels;
    Bavg = sum(B(:))/numpixels;

    Havg = sum(H(:))/numpixels;
    Savg = sum(S(:))/numpixels;

```

```

Vavg = sum(V(:))/numpixels;

outfile = ['HISTOGRAMS/PLANE/EXP1/image' int2str(j)];
eval(['save ', outfile variables]);
end

UNIFORM\_FV2

for i = 1:10

    infile = ['NORMALIZED_DATA/EXP_IMAGES_2/image' int2str(i)];
    [histA,histC1,histC2,histAf,histC1f,histC2f,histR,histG,histB,histH,histS,histV] = converti_
    outfile = ['HISTOGRAMS/UNIFORM/EXP1/image' int2str(i)];
    eval(['save ', outfile , histA histC1 histC2 histAf histC1f histC2f histR histG histB histH h
end

CONVERT\_IMAGES2

function [histA,histC1,histC2,histAf,histC1f,histC2f,histR,histG,histB,histH,histS,histV] = conver
eval(['load ', filename]);
load 'UNIFORM_HISTOGRAMS/db_min_max2';
%
% Calculate uniform histograms in the different planes
Astep = (Amax-Amin)/20;
C1step = (C1max-C1min)/20;
C2step = (C2max-C2min)/20;

Afstep = (Afmax-Afmin)/20;
C1fstep = (C1fmax-C1fmin)/20;
C2fstep = (C2fmax-C2fmin)/20;

Rstep = (Rmax-Rmin)/20;
Gstep = (Gmax-Gmin)/20;
Bstep = (Bmax-Bmin)/20;

Hstep = (Hmax-Hmin)/20;
Sstep = (Smax-Smin)/20;
Vstep = (Vmax-Vmin)/20;

Abins = Amin:Astep:Amax;
C1bins = C1min:C1step:C1max;
C2bins = C2min:C2step:C2max;

Afbins = Afmin:Afstep:Afmax;

```

```

C1fbins = C1fmin:C1fstep:C1fmax;
C2fbins = C2fmin:C2fstep:C2fmax;

Rbins = Rmin:Rstep:Rmax;
Gbins = Gmin:Gstep:Gmax;
Bbins = Bmin:Bstep:Bmax;

Hbins = Hmin:Hstep:Hmax;
Sbins = Smin:Sstep:Smax;
Vbins = Vmin:Vstep:Vmax;

histA = zeros(1,20);
histC1 = zeros(1,20);
histC2 = zeros(1,20);
histAf = zeros(1,20);
histC1f = zeros(1,20);
histC2f = zeros(1,20);
histR = zeros(1,20);
histG = zeros(1,20);
histB = zeros(1,20);
histH = zeros(1,20);
histS = zeros(1,20);
histV = zeros(1,20);

for i = 1:(size(A,1))^2
    for j = 1:20
        if ((A(i) >= Abins(j)) & ( A(i) < Abins(j+1)))
            histA(j) = histA(j)+1;
        end
        if ((C1(i) >= C1bins(j)) & ( C1(i) < C1bins(j+1)))
            histC1(j) = histC1(j)+1;
        end
        if ((C2(i) >= C2bins(j)) & ( C2(i) < C2bins(j+1)))
            histC2(j) = histC2(j)+1;
        end
        if ((Af(i) >= Afbins(j)) & ( Af(i) < Afbins(j+1)))
            histAf(j) = histAf(j)+1;
        end
        if ((C1f(i) >= C1fbins(j)) & ( C1f(i) < C1fbins(j+1)))
            histC1f(j) = histC1f(j)+1;
        end
        if ((C2f(i) >= C2fbins(j)) & ( C2f(i) < C2fbins(j+1)))
            histC2f(j) = histC2f(j)+1;
        end
        if ((R(i) >= Rbins(j)) & ( R(i) < Rbins(j+1)))
            histR(j) = histR(j)+1;
        end
        if ((G(i) >= Gbins(j)) & ( G(i) < Gbins(j+1)))
            histG(j) = histG(j)+1;
    end
end

```

```

        end
        if ((B(i) >= Bbins(j)) & ( B(i) < Bbins(j+1)))
histB(j) = histB(j)+1;
        end
        if ((H(i) >= Hbins(j)) & ( H(i) < Hbins(j+1)))
histH(j) = histH(j)+1;
        end
        if ((S(i) >= Sbins(j)) & ( S(i) < Sbins(j+1)))
histS(j) = histS(j)+1;
        end
        if ((V(i) >= Vbins(j)) & ( V(i) < Vbins(j+1)))
histV(j) = histV(j)+1;
        end
    end
end

```

EVAL_SIM_AVG2

```
for i = 1:10
```

```

infile = ['HISTOGRAMS/NONUNIFORM/EXP1/image' int2str(i)];
eval(['load ' infile]);
Aavg1=histA;
C1avg1=histC1;
C2avg1=histC2;
Afavg1=histAf;
C1favg1=histC1f;
C2favg1=histC2f;
Ravg1=histR;
Gavg1=histG;
Bavg1=histB;
Havg1=histH;
Savg1=histS;
Vavg1=histV;

```

```
for j = 1:10
```

```

infile = ['HISTOGRAMS/NONUNIFORM/EXP1/image' int2str(j)];
eval(['load ' infile]);
[rgbdist,hsvdist,ac1c2dist,ac1c2fdist] = euclidean_sim(Ravg1,Gavg1,Bavg1,Havg1,Savg1,Vavg1,Aavg1,
outputmatrix(i,j,1) = rgbdist;
outputmatrix(i,j,2) = hsvdist;
outputmatrix(i,j,3) = ac1c2dist;
outputmatrix(i,j,4) = ac1c2fdist;

```

```

    end

end
temp = max(max(outputmatrix(:,:,1)));
outputmatrix(:,:,1) = 10*(1-(outputmatrix(:,:,1)/((10/9)*temp)));
temp = max(max(outputmatrix(:,:,2)));
outputmatrix(:,:,2) = 10*(1-(outputmatrix(:,:,2)/((10/9)*temp)));
temp = max(max(outputmatrix(:,:,3)));
outputmatrix(:,:,3) = 10*(1-(outputmatrix(:,:,3)/((10/9)*temp)));
temp = max(max(outputmatrix(:,:,4)));
outputmatrix(:,:,4) = 10*(1-(outputmatrix(:,:,4)/((10/9)*temp)));

outfile = ['SIM_MEASURES/nonuni_norm'];
eval(['save ' outfile ' outputmatrix']);

EUCLIDEAN\_SIM

function [rgb_sim1,hsv_sim1,faug_sim1,faug_sim2] = euclidean_sim(histR_1,histG_1,histB_1,histH_1,)

r_diff=(histR_1-histR_2).^2;
g_diff=(histG_1-histG_2).^2;
b_diff=(histB_1-histB_2).^2;

rgb_sim1 = sum(sqrt(r_diff+g_diff+b_diff));

H_diff=(histH_1-histH_2).^2;
S_diff=(histS_1-histS_2).^2;
V_diff=(histV_1-histV_2).^2;

hsv_sim1= sum(sqrt(H_diff+S_diff+V_diff));

A_diff=(histA_1-histA_2).^2;
C1_diff=(histC1_1-histC1_2).^2;
C2_diff=(histC2_1-histC2_2).^2;

faug_sim1= sum(sqrt((A_diff)+(C1_diff)+(C2_diff)));

Af_diff=(histAf_1-histAf_2).^2;
C1f_diff=(histC1f_1-histC1f_2).^2;
C2f_diff=(histC2f_1-histC2f_2).^2;

faug_sim2= sum(sqrt((Af_diff)+(C1f_diff)+(C2f_diff)));

```

Appendix D. Test Images

Table D.1 Ten Test Images Used for Experiments

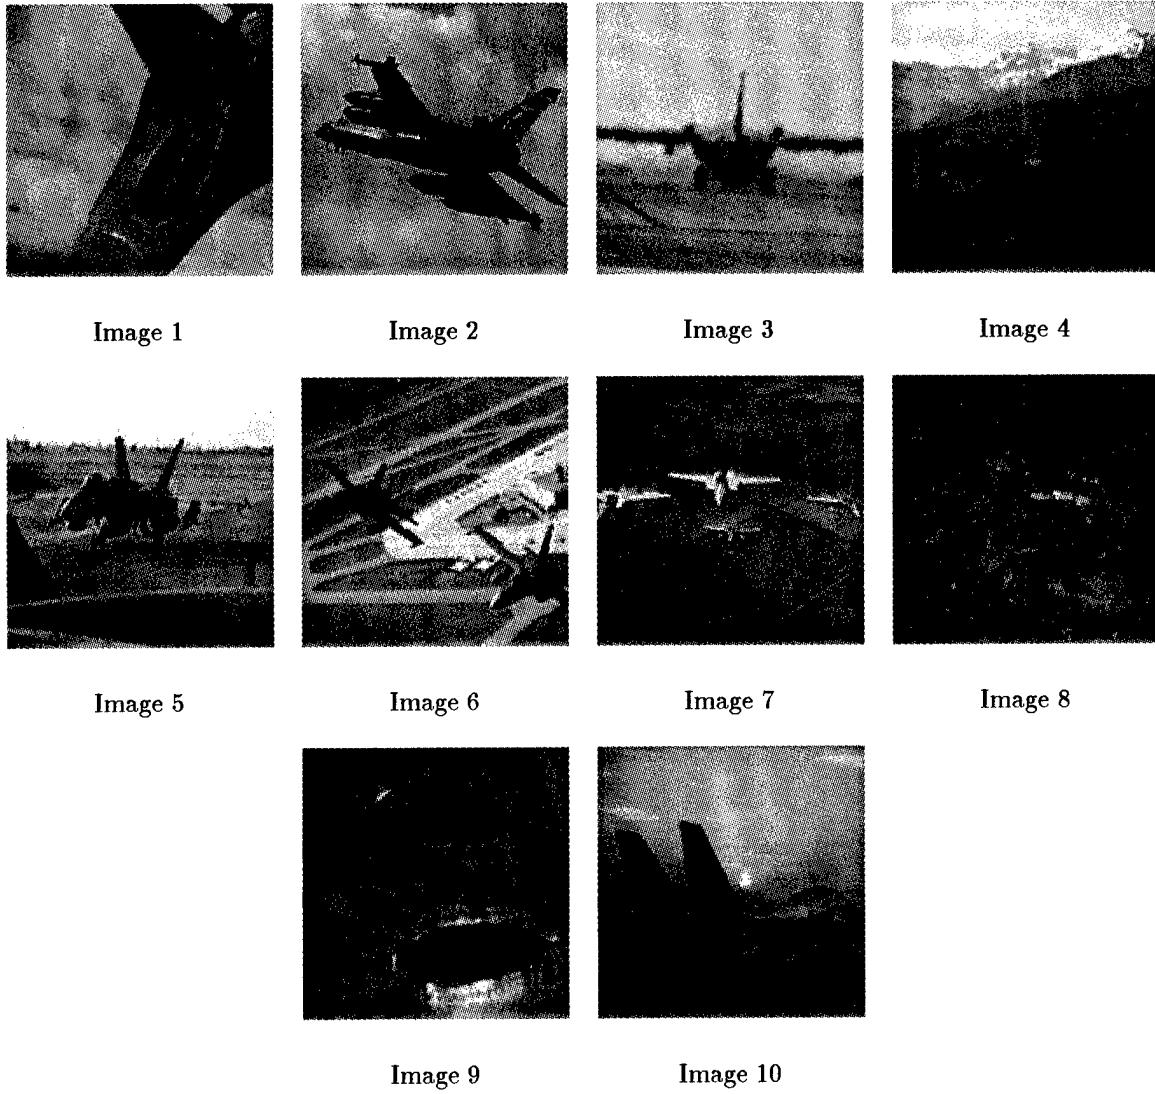


Table D.2 Images Used for Tutorial



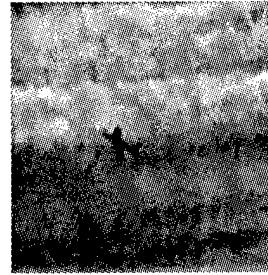
Tutorial Test Image



Example of Similarity



Example of Dissimilarity



Example of Medium Similarity

Bibliography

1. C. Cluff *et al.*, "Imaging: A vision of the future," Critical Technology Report C-10-1, Chantico Publishing Co., 1991.
2. A. K. Jain and A. Vailaya, "Image retrieval using color and shape," *Pattern Recognition*, vol. 29, pp. 1233–1243, January 1996.
3. B. Prabhakaran, *Multimedia Database Management Systems*. Boston/Dordrecht/London: Kluwer Academic Publishers, 1997.
4. W. K. Pratt, *Digital Image Processing*. New York: John Wiley & Sons, 1978.
5. C. Martin, *Perceptual Fidelity for Digital Color Imagery*. PhD thesis, Air Force Institute of Technology, Wright Patterson AFB, OH, 1996.
6. M. E. Ivie, "Future trends and requirements for databases in the pacs environment," *SPIE*, vol. 767, pp. 839–843, 1987.
7. A. J. Michaelis, "Images and the information superhighway," in *IS&T's 48th Annual Conference Proceedings*, pp. 19–21.
8. D. Forsyth *et al.*, "Searching for digital pictures," *Scientific American*, pp. 88–93, June 1997.
9. M. J. Swain and D. H. Ballard, "Color indexing," *International Journal of Computer Vision*, vol. 7, pp. 11–32, November 1991.
10. O. D. Faugeras, "Digital color image processing and psychophysics within the framework of a human visual model," Technical Report UTEC-CSc-77-029, University of Utah, Computer Science Department, University of Utah, Salt Lake City, Utah 84112, June 1976.
11. O. D. Faugeras, "Digital color image processing within the framework of a human visual model," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, pp. 380–393, August 1979.
12. W. Niblack *et al.*, "The qbic project: Querying images by content using color, texture, and shape," *SPIE*, vol. 1908, pp. 173–180, 1993.
13. Pentland *et al.*, "Photobook: Tools for content-based manipulation of image databases," *SPIE*, vol. 2185, pp. 34–43, 1994.
14. Bach *et al.*, "The virage image search engine: an open framework for image management," *SPIE*, vol. 2670, pp. 76–87, 1996.
15. R. L. Delanoy, "Toolkit for image mining: User-trainable search tools," *The Lincoln Laboratory Journal*, vol. 8, no. 2, pp. 145–160, 1995.
16. C. Faloutsos, *Searching Multimedia Databases by Content*. Boston/Dordrecht/London: Kluwer Academic Publishers, 1996.
17. Y. Niu *et al.*, "A study of image indexing techniques for multimedia database systems," Technical Report TR 95-19, University of Alberta, Laboratory for Database Systems Research, Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada T6G 2H1, July 1995.
18. D. Gavrilla, "R-tree index optimization," Technical Report CAR-TR-718, University of Maryland, Computer Vision Laboratory, Center for Automation Research, University of Maryland, College Park, MD 20742-3275, June 1994.
19. R. Jain and D. A. White, "Similarity indexing with the ss-tree," in *Proceedings of the International Conference on Data Engineering*, pp. 516–523, 1996.
20. G. Lu, "On image retrieval based on colour," *SPIE*, vol. 2670, pp. 310–319, 1996.

21. X. Wan and C.-C. J. Kuo, "Color distribution analysis and quantization for image retrieval," *SPIE*, vol. 2670, pp. 8–16, 1996.
22. S. J. Murthy *et al.*, "Similarity measures for image databases," *SPIE*, vol. 2420, pp. 58–62, 1995.
23. A. Gersho *et al.*, *Vector Quantization and Signal Compression*. Boston/Dordrecht/London: Kluwer Academic Publishers, 1992.
24. A. N. Netravali and B. G. Haskell, *Digital Pictures*. New York: Plenum Press, 1995.
25. B. A. Wandell, *Foundations of Vision*. Sunderland, Massachusetts: Sinauer Associates, Inc., 1995.
26. M. W. Schwarz *et al.*, "An experimental comparison of rgb, yiq, lab, hsv, and opponent color models," *ACM Transactions on Graphics*, vol. 6, pp. 123–158, April 1987.
27. C. A. Poynton, "Color space faq - frequently asked questions about color and gamma," *WWWeb* - <http://www.inforamp.net/poynton/Poynton-color.html>, February 1997.
28. T. A. Wilson, "Perceptual based image fusion with applications to hyperspectral image data," Master's thesis, Air Force Institute of Technology, Wright Patterson AFB, OH, 1994.
29. S. Zeki, *A Vision of the Brain*. Oxford: Blackwell Scientific Publications, 1993.
30. F. A. Maher, "A correlation of human and machine pattern discrimination," in *National Aerospace and Engineering Conference - NAECON*, pp. 260–264, 1970.
31. J. W. Barnes, *Statistical Analysis for Engineers and Scientists*. New York St. Louis San Francisco: McGraw-Hill, 1994.
32. D. Lane, "Hyperstat online," *WWWeb*-<http://www.ruf.rice.edu/lane/hyperstat/contents.html>, October 1997.

Vita

Second Lieutenant Chad A. Vander Meer was born in Billings, Montana on 3 June 1973. He graduated from Malad High School, Malad City, Idaho in 1991. After graduating from high school, Chad entered college at the University of Idaho. In 1996, he received a Bachelor of Science degree in Computer Engineering and was commissioned a second lieutenant. His first assignment was to the Air Force Institute of Technology, Wright-Patterson AFB, Ohio to pursue a Master of Science degree in Computer Engineering. Upon graduation, Lieutenant Vander Meer will be transferred to the USAFE Computer Systems Squadron, Ramstein Air Base, Germany.

Permanent address: P.O. Box 565
Parma, ID 83660

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</p>			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	December 1997	Master's Thesis	
4. TITLE AND SUBTITLE A PERFORMANCE ANALYSIS OF THE FAUGERAS COLOR SPACE AS A COMPONENT OF COLOR HISTOGRAM-BASED IMAGE RETRIEVAL			5. FUNDING NUMBERS
6. AUTHOR(S) Chad A. Vander Meer, 2d Lt, USAF			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology, WPAFB, OH 45433-7765			8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GCS/ENG/97D-18
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) RL/IRDS 32 Hangar Road Rome, NY 13441-4114			10. SPONSORING/MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The use of color histograms for image retrieval from databases has been implemented in many variations since the original work of Ballard and Swain. Selecting the appropriate color space for similarity comparisons is an important part of a color histogram technique. This paper serves to introduce and evaluate the performance of a color space developed by O.D. Faugeras through the use of color histograms. Performance is evaluated by correlating the similarity results obtained from various color feature vector techniques (including color histogramming) to those gathered through a human perceptual test. The perceptual test required 36 human subjects to evaluate the similarity of 10 military aircraft images. The same 10 images were also compared via the color feature vector techniques. The results obtained for the Faugeras color space are compared against those of the Red, Green, Blue (RGB) and Hue, Saturation, Value (HSV) color spaces. While the correlation results for the Faugeras color space were unexpected and unfavorable, a Pearson correlation coefficient of 0.91 was obtained for the HSV space suggesting that HSV is an excellent color space for judging color image similarity. A discussion of the Faugeras space's performance and future research directions are presented at the conclusion of the paper.			
14. SUBJECT TERMS Color Histograms, Color Spaces, Color Space Quantization, Human Visual System, Content-Based Image Retrieval		15. NUMBER OF PAGES 96 16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL